Accelerated and Inexact Soft-Impute for Large-Scale Matrix and Tensor Completion

Quanming Yao James T. Kwok
Department of Computer Science and Engineering
Hong Kong University of Science and Technology
Clear Water Bay, Hong Kong
{qyaoaa, jamesk}@cse.ust.hk

arXiv:1703.05487v1 [cs.NA] 16 Mar 2017

Abstract

Matrix and tensor completion aim to recover a low-rank matrix / tensor from limited observations and have been commonly used in applications such as recommender systems and multi-relational data mining. A state-of-the-art matrix completion algorithm is Soft-Impute, which exploits the special "sparse plus low-rank" structure of the matrix iterates to allow efficient SVD in each iteration. Though Soft-Impute is a proximal algorithm, it is generally believed that acceleration destroys the special structure and is thus not useful. In this paper, we show that Soft-Impute can indeed be accelerated without comprising this structure. To further reduce the iteration time complexity, we propose an approximate singular value thresholding scheme based on the power method. Theoretical analysis shows that the proposed algorithm still enjoys the fast $O(1/T^2)$ convergence rate of accelerated proximal algorithms. We further extend the proposed algorithm to tensor completion with the scaled latent nuclear norm regularizer. We show that a similar "sparse plus low-rank" structure also exists, leading to low iteration complexity and fast $O(1/T^2)$ convergence rate. Extensive experiments demonstrate that the proposed algorithm is much faster than Soft-Impute and other state-of-the-art matrix and tensor completion algorithms.

1 Introduction

Matrices are common place in data mining applications. For example, in recommender systems, the ratings data can be represented as a sparsely observed user-item matrix [Adomavicius and Tuzhilin, 2005; Koren, 2008]. In social networks, user interactions can be modeled by an adjacency matrix [Kim and Leskovec, 2011; Chiang *et al.*, 2014]. Matrices also appear in applications such as image processing [Liu *et al.*, 2013; Haeffele *et al.*, 2014; Yao and Kwok, 2015b] and question answering [Zhao *et al.*, 2015].

Due to limited feedback from users, these matrices are usually not fully observed. For example, users may only give opinions on very few items in a recommender system. As the rows/columns are usually related to each other, the low-rank matrix assumption is particularly useful to capture such relatedness, and low-rank matrix completion has become a powerful tool to predict missing values in these matrices. Sound recovery guarantee [Candès and Recht, 2009] and good empirical performance [Koren, 2008] have been obtained.

However, directly minimizing the matrix rank is NP-hard [Recht *et al.*, 2010]. To alleviate this problem, the nuclear norm (which is the sum of singular values) is often used instead. It is known that the nuclear norm is the tightest convex lower bound of the rank [Recht *et al.*, 2010]. Specifically, consider an $m \times n$ matrix O (without loss of generality, we assume that $m \geq n$), with positions of the observed entries indicated by $\Omega \in \{0,1\}^{m\times n}$, where $\Omega_{ij} = 1$ if O_{ij} is observed, and 0 otherwise. The matrix completion problem can be formulated as

$$\min_X \frac{1}{2}\|P_\Omega(X - O)\|_F^2 + \lambda\|X\|_*, \tag{1}$$

where $[P_\Omega(A)]_{ij} = A_{ij}$ if $\Omega_{ij} = 1$, and 0 otherwise; and $\|\cdot\|_*$ is the nuclear norm. Though the nuclear norm is only a surrogate of the matrix rank, there are theoretical guarantees that the underlying matrix can be exactly recovered [Candès and Recht, 2009].

Computationally, though the nuclear norm is nonsmooth, problem (1) can be solved by various optimization tools. An early attempt is based on reformulating (1) as a semidefinite program (SDP) [Candès and Recht, 2009]. However, SDP solvers have large time and space complexities, and are only suitable for small data sets. For large-scale matrix completion, singular value thresholding (SVT) algorithm [Cai *et al.*, 2010] pioneered the use of first-order methods . However, a singular value decomposition (SVD) is required in each SVT iteration. This takes $O(mn^2)$ time and can be computationally expensive. In [Toh and Yun, 2010], this is reduced to a partial SVD by computing only the leading singular values/vectors using PROPACK (a variant of the Lanczos algorithm) [Larsen, 1998]. Another major breakthrough is made by the Soft-Impute algorithm [Mazumder *et al.*, 2010], which utilizes a special "sparse plus low-rank" structure associated with the SVT to efficiently compute the SVD. Empirically, this allows Soft-Impute to perform matrix completion on the entire *Netflix* data set. The SVT algorithm

can also be viewed as a proximal algorithm [Tibshirani, 2010]. Hence, it converges with a $O(1/T)$ rate, where T is the number of iterations [Beck and Teboulle, 2009; Nesterov, 2013]. Later, this is further "accelerated", and the convergence rate is improved to $O(1/T^2)$ [Ji and Ye, 2009; Toh and Yun, 2010]. However, Tibshirani [Tibshirani, 2010] suggested that this is not useful, as the special "sparse plus low-rank" structure crucial to the efficiency of Soft-Impute no longer exist. In other words, the gain in convergence rate is more than compensated by the increase in iteration time complexity.

In this paper, we show that accelerating Soft-Impute is indeed possible while still preserving the "sparse plus low-rank" structure. To further reduce the iteration time complexity, instead of computing SVT exactly using PROPACK [Toh and Yun, 2010; Mazumder *et al.*, 2010], we propose an approximate SVT scheme based on the power method [Halko *et al.*, 2011]. Though the SVT obtained in each iteration is only approximate, we show that convergence can still be as fast as performing exact SVT. Hence, the resultant algorithm has low iteration complexity and fast $O(1/T^2)$ convergence rate. To further boost performance, we propose a post-processing procedure by extending [Mazumder *et al.*, 2010], and can handle any smooth convex loss function.

Besides matrices, tensors have also been commonly used to describe the linear and multilinear relationships in the data [Kolda and Bader, 2009; Liu *et al.*, 2013; Kazienko *et al.*, 2011; Shin *et al.*, 2017]. For example, in remote sensing applications, a hyperspectral image with multiple bands can be naturally represented as a 3-dimensional tensor. A multidimensional social network can also be modeled as a 3-dimensional tensor, where the third mode may represent different type of relations. Higher-dimensional tensors are also useful. For example, a multi-mode social network (such as the DBLP network) with heterogeneous actors (papers, authors, terms and venues) can be represented by a 4-order tensor, and a relation can connect these four kinds of entities [Tang *et al.*, 2008]. Analogous to matrix completion, tensor completion attempts to recover a low-rank tensor that best approximates a partially observed data tensor [Tomioka *et al.*, 2010; Liu *et al.*, 2013]. For example, in hyperspectral imaging, as some bands may be partially missing due to sensor problems, tensor completion can be used to inpaint the incomplete image [Xu *et al.*, 2013].

Analogous to matrix completion, tensor completion can also be solved by convex optimization algorithms. However, multiple expensive SVDs on large dense matrices are required [Liu *et al.*, 2013; Tomioka *et al.*, 2010]. To alleviate this problem, we demonstrate that a similar "sparse plus low-rank" structure also exists when the scaled latent nuclear norm [Tomioka *et al.*, 2010; Wimalawarne *et al.*, 2014] is used as the regularizer. We extend the proposed matrix-based algorithm to this tensor scenario. The resulting algorithm has low iteration cost and fast $O(1/T^2)$ convergence rate. Experiments on matrix/tensor completion problems with both synthetic and real-world data sets show that the proposed algorithm outperforms state-of-the-art algorithms.

Preliminary results of this paper have been reported in a shorter conference version [Yao and Kwok, 2015a]. While only the square loss is used in [Yao and Kwok, 2015a], here we consider more general smooth convex loss functions. Moreover, we extend the proposed algorithm to tensor completion. Besides, post-processing is proposed to boost the recovery performance for matrix/tensor completion.

The rest of the paper is organized as follows. Section 2 provides a brief review on the related work. The proposed accelerated inexact Soft-Impute algorithm is described in Section 3, and its extension to tensor completion in Section 4. Experimental results are presented in Section 5, and the last section gives some concluding remarks. All the proofs are in the appendix.

Notation

In the sequel, the transpose of vector/ matrix is denoted by the superscript $\cdot^\top$, and tensors are denoted by boldface Euler. For a vector x, $\|x\|_1 = \sum_i |x_i|$ is its ℓ_1-norm, and $\|x\| = \sqrt{\sum_i x_i^2}$ its ℓ_2-norm. For a matrix X, $\sigma_1(X) \geq \sigma_2(X) \geq \ldots \sigma_m(X)$ are its singular values, $\mathrm{tr}(X) = \sum_i X_{ii}$ is its trace, $\|X\|_1 = \sum_{i,j} |X_{ij}|$, $\|X\|_\infty$ is its maximum singular value, and $\|X\|_F = \mathrm{tr}(X^\top X)$ the Frobenius norm, $\|X\|_* = \sum_i \sigma_i(X)$ the nuclear norm, and $\mathrm{span}(X)$ is the column span of X. Moreover, I denotes the identity matrix.

For tensors, we follow the notations in [Kolda and Bader, 2009]. For a D-order tensor $\mathcal{X} \in \mathbb{R}^{I_1 \times I_2 \times \cdots \times I_D}$, its $(i_1, i_2, \ldots, i_D)$th entry is $x_{i_1 i_2 \ldots i_D}$. Let $I_{D\backslash d} = \prod_{j=1, j\neq d}^{D} I_j$, the mode-$d$ matricizations $\mathcal{X}_{\langle d \rangle}$ of $\mathcal{X}$ is a $I_d \times I_{D\backslash d}$ matrix with $(\mathcal{X}_{\langle d \rangle})_{i_d j} = x_{i_1 i_2 \cdots i_D}$, and $j = 1 + \sum_{l=1, l\neq d}^{D} (i_l - 1) \prod_{m=1, m\neq d}^{l-1} I_m$. Given a matrix A, its mode-d tensorization $A^{\langle d \rangle}$ is a tensor $\mathcal{X}$ with elements $x_{i_1 i_2 \cdots i_D} = a_{i_d j}$, and j is as defined above. The inner product of two tensors $\mathcal{X}$ and $\mathcal{Y}$ is $\langle \mathcal{X}, \mathcal{Y} \rangle = \sum_{i_1=1}^{I_1} \cdots \sum_{i_D=1}^{I_D} x_{i_1 i_2 \ldots i_D} y_{i_1 i_2 \ldots i_D}$, and the Frobenius norm of $\mathcal{X}$ is $\|\mathcal{X}\|_F = \sqrt{\langle \mathcal{X}, \mathcal{X} \rangle}$.

For a convex but nonsmooth function f, the subgradient is $g \in \partial f(x)$ where $\partial f(x) = \{u : f(y) \geq f(x) + u^\top (y - x), \forall y\}$ is its subdifferential. When f is differentiable, we use ∇f for its gradient.

2 Related Work

2.1 Proximal Algorithms

Consider minimizing composite functions of the form:

$$F(x) \equiv f(x) + g(x), \tag{2}$$

where f, g are convex, and f is smooth but g is possibly non-smooth. The proximal algorithm [Parikh and Boyd, 2014] generates a sequence of estimates $\{x_t\}$ as

$$x_{t+1} = \mathrm{prox}_{\mu g}(z_t) \equiv \arg\min_x \frac{1}{2}\|x - z_t\|_2^2 + \mu g(x),$$

where

$$z_t = x_t - \mu \nabla f(x_t), \tag{3}$$

and $\mathrm{prox}_{\mu g}(\cdot)$ is the proximal operator. When f is ρ-Lipschitz smooth (i.e., $\|\nabla f(x_1) - \nabla f(x_2)\| \leq \rho \|x_1 - x_2\|$) and a fixed stepsize

$$\mu \leq 1/\rho \tag{4}$$

is used, the proximal algorithm converges to the optimal solution with a rate of $O(1/T)$, where T is the number of iterations [Combettes and Wajs, 2005; Parikh and Boyd, 2014]. By replacing the update in (3) with

$$y_t = (1+\theta_t)x_t - \theta_t x_{t-1}, \quad (5)$$
$$z_t = y_t - \mu\nabla f(y_t), \quad (6)$$

where $\theta_{t+1} = \frac{t-1}{t+2}$, it can be accelerated to a convergence rate of $O(1/T^2)$ [Beck and Teboulle, 2009; Nesterov, 2013]. This is also known to be the best possible rate for problems of the form in (2) [Nesterov, 2013].

Algorithm 1 Accelerated proximal gradient (APG) algorithm.

initialize $x_0 = x_1 = 0$ and stepsize $\mu \in (0, 1/\rho]$;
for $t = 1, 2, \cdots$ **do**
  $\theta_t = \frac{t-1}{t+2}$;
  $y_t = (1+\theta_t)x_t - \theta_t x_{t-1}$;
  $z_t = y_t - \mu\nabla f(y_t)$;
  $x_{t+1} = \text{prox}_{\mu g}(z_t)$;
end for
return x_{t+1}.

Often, g is "simple" in the sense that $\text{prox}_{\mu g}(\cdot)$ can be easily obtained. However, in more complicated problems such as overlapping group lasso [Jacob *et al.*, 2009], $\text{prox}_{\mu g}(\cdot)$ may be expensive to compute. To alleviate this problem, inexact proximal algorithm is proposed which allows two types of errors in standard/accelerated proximal algorithms [Schmidt *et al.*, 2011]: (i) an error e_t in computing $\nabla f(\cdot)$, and (ii) an error ε_t in the proximal step, i.e.,

$$h_{\mu g}(x_{t+1}; z_t) \le \varepsilon_t + h_{\mu g}(\text{prox}_{\mu g}(z_t); z_t), \quad (7)$$

where

$$h_{\mu g}(x; z_t) \equiv \frac{1}{2}\|x - z_t\|^2 + \mu g(x) \quad (8)$$

is the proximal step's objective. Let the dual problem of $\min_x h_{\mu g}(x; z_t)$ be $\max_w \mathcal{D}_{\mu g}(w)$. Note that ε_t is upper-bounded by the duality gap $\vartheta_t \equiv h_{\mu g}(x_{t+1}; z_t) - \mathcal{D}(w_{t+1})$ where w_{t+1} is the corresponding dual variable of x_{t+1}. Thus, (7) can be ensured by monitoring ϑ_t. The following Proposition shows that by decreasing e_t and ε_t sufficiently fast, the convergence rate remains at $O(1/T^2)$.

Proposition 2.1 ([Schmidt *et al.*, 2011]). *If $\|e_t\|$ and $\sqrt{\varepsilon_t}$ decrease as $O(1/t^{2+\delta})$ for some $\delta > 0$, the inexact accelerated proximal gradient algorithm converges with a rate of $O(1/T^2)$.*

In the sequel, as our focus is on matrix completion, the variable x in (2) will be a matrix X.

2.2 Soft-Impute

Soft-Impute [Mazumder *et al.*, 2010] is a state-of-the-art algorithm for matrix completion (Algorithm 2). At iteration t, let the current iterate be X_t. The missing values in O are filled in as

$$Z_t = P_\Omega(O) + P_{\Omega^\perp}(X_t) = P_\Omega(O - X_t) + X_t, \quad (9)$$

where $\Omega^\perp$ is the complement of Ω (i.e., $\Omega^\perp_{ij} = 1 - \Omega_{ij}$). The next estimate X_{t+1} is then generated by the singular value thresholding (SVT) operator [Cai *et al.*, 2010]

$$X_{t+1} = \text{SVT}_l(Z_t) \equiv \arg\min_X \frac{1}{2}\|X - Z_t\|_F^2 + \lambda\|X\|_*, \quad (10)$$

which can be computed as follows.

Lemma 2.2 ([Cai *et al.*, 2010]). *Let the SVD of a matrix Z_t be $U\Sigma V^\top$. Then, $\text{SVT}_l(Z_t) \equiv U(\Sigma - \lambda I)_+ V^\top$ where $[(A)_+]_{ij} = \max(A_{ij}, 0)$.*

Algorithm 2 Soft-Impute.

Require: Partially observed matrix O, parameter λ;
initialize $X_1 = 0$;
for $t = 1, 2, \cdots$ **do**
  $Z_t = P_\Omega(O) + P_{\Omega^\perp}(X_t)$;
  $X_{t+1} = \text{SVT}_\lambda(Z_t)$;
end for
return X_{t+1}.

Let $\bar{k}_t$ be the number of singular values in Z_t that are larger than λ. From Lemma 2.2, a rank-k_t SVD, where $k_t \ge \bar{k}_t$, is sufficient for computing X_{t+1} in (10). In [Mazumder *et al.*, 2010], this rank-k_t SVD is obtained by the PROPACK algorithm [Larsen, 1998].

To make Soft-Impute efficient, an important observation in [Mazumder *et al.*, 2010] is that Z_t in (9) has a special "sparse plus low-rank" structure, namely that $P_\Omega(O - X_t)$ is sparse and X_t is low-rank. The most expensive steps in computing the SVD are matrix-vector multiplications of the form Zu and $v^\top Z$, where $u \in \mathbb{R}^n$ and $v \in \mathbb{R}^m$. Let the rank of X_t be r_t, and its SVD be $U_t\Sigma_t{V_t}^\top$. $Z_t v$ can be computed as

$$Z_t v = P_\Omega(O - X_t)v + U_t\Sigma_t({V_t}^\top v). \quad (11)$$

Constructing $P_\Omega(O - X_t)$ takes $O(r_t\|\Omega\|_1)$ time, while computing the products $P_\Omega(O - X_t)u$ and $U_t\Sigma_t({V_t}^\top u)$ take $O(\|\Omega\|_1)$ and $O(mr_t)$ time, respectively. Similarly, $u^\top Z_t$ can be computed as $u^\top P_\Omega(O - X_t) + (u^\top U_t)\Sigma_t{V_t}^\top$. Thus, to obtain the rank-k SVD of Z_t, Soft-Impute needs only

$$O(k_t\|\Omega\|_1 + r_t k_t m) \quad (12)$$

time, and one iteration costs

$$O((r_t + k_t)\|\Omega\|_1 + r_t k_t m) \quad (13)$$

time. Since the solution is low-rank, $k_t, r_t \ll m$, and (13) is much faster than the $O(mnk_t)$ time for direct rank-k_t SVD.

3 Accelerated Inexact Soft-Impute

In this section, we describe the proposed matrix completion algorithm. Tibshirani [Tibshirani, 2010] suggested that acceleration is not useful, as it destroys the essential "sparse plus low-rank" structure. However, we will show that it can indeed be preserved with acceleration. We also show that further speedup can be achieved by using approximate SVT.

3.1 Soft-Impute as a Proximal Algorithm

In (1), let

$$f(X) = \frac{1}{2}\|P_\Omega(X - O)\|_F^2 = \sum_{(i,j)\in\Omega} \ell(X_{ij}, O_{ij}), \quad (14)$$

where ℓ is the loss function, and $g(X) = \lambda\|X\|_*$. The proximal step in the (unaccelerated) proximal algorithm is

$$X_{t+1} = \text{prox}_{\mu g}(Z_t) \equiv \arg\min_x \frac{1}{2}\|X - Z_t\|_F^2 + \mu\lambda\|X\|_*,$$

where $Z_t = X_t - \mu P_\Omega(X_t - O)$. Note that the square loss $\ell(X_{ij}, O_{ij}) \equiv \frac{1}{2}(X_{ij} - O_{ij})^2$ in (1) is 1-Lipschitz smooth. The following shows that f in (14) is also 1-Lipschitz smooth. The proof can be found in Appendix A.1.

Proposition 3.1. *If ℓ is ρ-Lipschitz smooth, f in (14) is also ρ-Lipschitz smooth.*

From (4), one can thus simply set $\mu = 1$ for (1). We then have $X_{t+1} = \text{prox}_g(Z_t) = \text{SVT}_\lambda(Z_t)$ which is the same as (10). Hence, interestingly, Soft-Impute is a proximal algorithm [Tibshirani, 2010], and thus converges at a rate of $O(1/T)$ [Mazumder *et al.*, 2010].

3.2 Accelerating Soft-Impute

Since Soft-Impute is a proximal algorithm, it is natural to accelerate it (Section 2.1). In this section, we show that the "sparse plus low-rank" structure can also be preserved.

To accelerate Soft-Impute, recall from (5) and (6) that we have to compute

$$\begin{aligned}\text{prox}_g(\breve{Z}_t) &= \text{SVT}_\lambda(\breve{Z}_t) \\ &= \arg\min_X \frac{1}{2}\|X - \breve{Z}_t\|_F^2 + \lambda\|X\|_*, \quad (15)\end{aligned}$$

where

$$\begin{aligned} Y_t &= (1+\theta_t)X_t - \theta_t X_{t-1}, \\ \breve{Z}_t &= P_\Omega(O - Y_t) + (1+\theta_t)X_t - \theta_t X_{t-1}. \quad (16)\end{aligned}$$

Assume that X_t and X_{t-1} have ranks r_t and r_{t-1}, and their SVDs are $U_t\Sigma_t V_t^\top$ and $U_{t-1}\Sigma_{t-1}V_{t-1}^\top$, respectively. Similar to (11), for any $v \in \mathbb{R}^n$, we have $\breve{Z}_t v = P_\Omega(O - Y_t)v + (1+\theta_t)U_t\Sigma_t(V_t^\top v) - \theta_t U_{t-1}\Sigma_{t-1}(V_{t-1}^\top v)$. The first term takes $O(\|\Omega\|_1)$ time while the last two terms take $O((r_{t-1}+r_t)m)$ time, thus a total of $O(\|\Omega\|_1 + (r_{t-1}+r_t)m)$ time. Similarly, for any $u \in \mathbb{R}^m$, $u^\top \breve{Z}_t$ takes $O(\|\Omega\|_1 + (r_{t-1}+r_t)m)$ time. The rank-k_t SVD of $\breve{Z}_t$ can be obtained using PROPACK in

$$O(k_t\|\Omega\|_1 + (r_{t-1}+r_t)k_t m) \quad (17)$$

time. As the target matrix is low-rank, r_{t-1} and r_t are much smaller than n. Hence, (17) is much faster than the $O(mnk_t)$ time required for a direct rank-k_t SVD.

The accelerated algorithm has a slightly higher iteration complexity than the unaccelerated one in (12). However, this is more than compensated by improvement in the convergence rate (from $O(1/T)$ to $O(1/T^2)$), as will be empirically demonstrated in Section 5.1.

3.3 Approximating the SVT

Though acceleration preserves the "sparse plus low-rank" structure, the proposed algorithm (and Soft-Impute) can still be computationally expensive as the SVT in each iteration uses exact SVD. In this section, we show that further speedup is possible by using inexact SVD.

As SVT in (10) can be seen as a proximal step, one might want to perform inexact SVT by monitoring the duality gap as in Section 2.1. It can be shown that the dual of (15) is

$$\max_{\|W\|_\infty \le 1} \text{tr}(W^\top \breve{Z}_t) - \frac{\lambda}{2}\|W\|_F^2, \quad (18)$$

where $W \in \mathbb{R}^{m\times n}$ is the dual variable.

Proposition 3.2 ([Parikh and Boyd, 2014])**.** *Let the SVD of matrix $\breve{Z}_t$ be $U\Sigma V^\top$. The optimal solution of* (18) *is $W_* = U\min(\Sigma, \lambda I)V^\top$, where $[\min(A,B)]_{ij} = \min(A_{ij}, B_{ij})$.*

Proposition 3.2 shows that a full SVD is required. This takes $O(m^2 n)$ time and is even more expensive than directly using SVT ($O(mnk_t)$ time). Instead, the proposed approximation is motivated by the following Proposition. The proof can be found in Appendix A.2.

Proposition 3.3. *Let $\breve{k}_t$ be the number of singular values in $\breve{Z}_t$ larger than λ, and $Q \in \mathbb{R}^{m\times k_t}$, where $k_t \ge \breve{k}_t$, be orthogonal and contains the subspace spanned by the top $\breve{k}_t$ left singular vectors of $\breve{Z}_t$. Then, $SVT_\lambda(\breve{Z}_t) = QSVT_\lambda(Q^\top \breve{Z}_t)$.*

Since a low-rank solution is desired, k_t can be much smaller than m [Mazumder *et al.*, 2010]. Thus, once we identify the span of $\breve{Z}_t$'s top left singular vectors, we only need to perform SVT on the much smaller $Q^\top \breve{Z}_t \in \mathbb{R}^{k_t\times n}$ (instead of $\breve{Z} \in \mathbb{R}^{m\times n}$). The question is how to find Q. We adopt the power method (Algorithm 3) [Halko *et al.*, 2011], which is more efficient than PROPACK [Wu and Simon, 2000]. Matrix R_t in Algorithm 3 is for warm-start.

Algorithm 3 PowerMethod($\breve{Z}_t, R_t, J$).

Require: $\breve{Z}_t \in \mathbb{R}^{m\times n}$, $R_t \in \mathbb{R}^{n\times k_t}$, and the number of iterations J;

1: initialize $Q_0 = \text{QR}(\breve{Z}_t R_t)$; // QR$(\cdot)$ is QR factorization

2: **for** $j = 1, 2, \ldots, J$ **do**

3: $\quad Q_j = \text{QR}(\breve{Z}_t(\breve{Z}_t^\top Q_{j-1}))$;

4: **end for**

5: **return** Q_J.

Algorithm 4 shows the approximate SVT procedure. Step 1 approximates the top k_t left singular vectors of $\breve{Z}_t$ with Q. In steps 2 to 5, a much smaller and less expensive (exact) SVT is performed on $Q^\top \breve{Z}_t$. Finally, $\text{SVT}_\lambda(\breve{Z}_t)$ is recovered as $\tilde{X} = (QU)\Sigma V^\top$ using Proposition 3.3.

Algorithm 4 Approximating the SVT of $\breve{Z}_t$: approx-SVT$(\breve{Z}_t, R_t, \lambda, J)$

Require: $\breve{Z}_t \in \mathbb{R}^{m\times n}$, $R_t \in \mathbb{R}^{n\times k_t}$ and $\lambda \geq 0$;
1: $Q = \text{PowerMethod}(\breve{Z}_t, R_t, J)$;
2: $[U, \Sigma, V] = \text{SVD}(Q^\top \breve{Z}_t)$;
3: $U = \{u_i \mid \sigma_i > \lambda\}$;
4: $V = \{v_i \mid \sigma_i > \lambda\}$;
5: $\Sigma = (\Sigma - \lambda I)_+$;
6: **return** QU, Σ and V. // $\tilde{X} = (QU)\Sigma V$

3.4 The Proposed Algorithm

We extend problem (1) by allowing the loss ℓ to be ρ-Lipschitz smooth (e.g., logistic loss and squared hinge loss):

$$\min_X F(X) \equiv \sum_{(i,j)\in\Omega} \ell(X_{ij}, O_{ij}) + \lambda\|X\|_*. \tag{19}$$

Using (6),

$$\breve{Z}_t = Y_t - \mu\nabla f(Y_t) = Y_t - \mu S_t,$$

where S_t is a sparse matrix with

$$[S_t]_{ij} = \begin{cases} \frac{d\ell((Y_t)_{ij}, O_{ij})}{d(Y_t)_{ij}} & \text{if } (i,j)\in\Omega \\ 0 & \text{otherwise} \end{cases}. \tag{20}$$

Using Proposition 3.1 and (4), the stepsize μ can be set as $1/\rho$. The whole procedure is shown in Algorithm 5. The core steps are 6–8, which performs approximate SVT. As in [Hsieh and Olsen, 2014], R_t is warm-started as QR$([V_t, V_{t-1}])$ at step 7. Moroever, as in [O'Donoghue and Candès, 2012; Nesterov, 2013], we restart the algorithm if $F(X)$ starts to increase (step 10). For further speedup, λ is dynamically reduced (step 3) by a continuation strategy [Toh and Yun, 2010; Mazumder *et al.*, 2010].

Algorithm 5 Accelerated Inexact Soft-Impute (AIS-Impute).

Require: partially observed matrix O, parameter λ.
1: initialize $c = 1$, $X_0 = X_1 = 0$, stepsize $\mu = 1/\rho$, $\hat{\lambda} > \lambda$ and $\nu \in (0,1)$;
2: **for** $t = 1, 2, \ldots, T$ **do**
3: $\quad\lambda_t = (\hat{\lambda} - \lambda)\nu^{t-1} + \lambda$;
4: $\quad Y_t = X_t + \theta_t(X_t - X_{t-1})$, where $\theta_t = \frac{c-1}{c+2}$;
5: $\quad\breve{Z}_t = Y_t - \mu S_t$, with S_t in (20);
6: $\quad V_{t-1} = V_{t-1} - V_t(V_t{}^\top V_{t-1})$, remove zero columns;
7: $\quad R_t = \text{QR}([V_t, V_{t-1}])$;
8: $\quad[U_{t+1}, \Sigma_{t+1}, V_{t+1}] = \text{approx-SVT}\left(\breve{Z}_t, R_t, \mu\lambda_t, J\right)$; // $X_{t+1} = U_{t+1}\Sigma_{t+1}V_{t+1}^\top$
9: $\quad$**if** $F(X_{t+1}) > F(X_t)$ **then**
10: $\quad\quad c = 1$;
11: $\quad$**else**
12: $\quad\quad c = c + 1$;
13: $\quad$**end if**
14: **end for**
15: **return** U_{T+1}, Σ_{T+1} and V_{T+1}.

3.5 Convergence and Time Complexity

In the following, we will show that the proposed algorithm has a convergence rate of $O(1/T^2)$. Let $X_{t+1} = U_{t+1}\Sigma_{t+1}V_{t+1}^\top$ be the output of approx-SVT at step 8. Since it only approximates $\text{SVT}_{\mu\lambda}(\breve{Z}_t)$, there is a difference (ε_t in (7)) between the proximal objectives $h_{\mu\lambda\|\cdot\|_*}(X_{t+1}; \breve{Z}_t)$ and $h_{\mu\lambda\|\cdot\|_*}(\text{SVT}_{\mu\lambda}(\breve{Z}_t); \breve{Z}_t)$ after performing step 8, where $h_{\mu\lambda\|\cdot\|_*}(\cdot;\cdot)$ is as defined in (8). The following shows that ε_t decreases at a linear rate. The proof can be found in Appendix A.3.

Proposition 3.4. *Assume that (i) $k_t \geq \breve{k}_t$ for all t and $J = t$;*[1] *(ii) $\{F(X_t)\}$ is upper-bounded. Then ε_t decreases to zero linearly.*

Using Propositions 2.1 and 3.4, convergence of the proposed algorithm is provided by the following Theorem. The proof can be found in Appendix A.4.

Theorem 3.5. *The sequence $\{X_t\}$ generated from Algorithm 5 converges to the optimal solution with a $O(1/T^2)$ rate.*

The basic operations in the power method are multiplications of the form $\breve{Z}_t u$ and $v^\top \breve{Z}_t$. The tricks in Section 3.2 can again be used for acceleration, and computing the approximate SVT using Algorithm 4 takes only

$$O(k_t\|\Omega\|_1 + (r_{t-1} + r_t)k_t m) \tag{21}$$

time. This is slightly more expensive than (12), the time for performing exact SVD in Soft-Impute. However, Soft-Impute is not accelerated and has slower convergence than Algorithm 5 (Theorem 3.5). The complexity in (21) is also the same as (17). However, as will be demonstrated in Section 5.1, approximate SVT is empirically much faster. Table 1 shows a breakdown of the iteration time complexity of Algorithm 5. As can be seen, it is only slightly more expensive than (13) for Soft-Impute.

Table 1: Iteration time complexity of Algorithm 5.

steps	complexity
5 (construct S_t)	$O(r_t\|\Omega\|_1)$
6,7 (warm-start)	$O(nk_t^2)$
8 (approximate SVT)	$O(k_t\|\Omega\|_1 + (r_{t-1} + r_t)\, k_t m)$
total	$O((r_t + k_t)\|\Omega\|_1 + (r_{t-1} + r_t + k_t)k_t m)$

Table 2 compares Algorithm 5 with some existing algorithms that will be empirically compared in Section 5.2. Overall, Algorithm 5 enjoys fast convergence and low iteration complexity.

3.6 Post-Processing

Recall that the nuclear norm penalizes all singular values equally. This may over-penalize the more important leading singular values. To alleviate this problem, we post-process the solution as in [Mazumder *et al.*, 2010]. However, only the

[1] In practice, we simply set $J = 3$ as in [Hsieh and Olsen, 2014].

Table 2: Comparison of AIS-Impute (Algorithm 5) and other algorithms. The algorithms in active subspace selection, TR and boost involve solving some optimization subproblems iteratively, and T_a is the number of iterations used. Moreover, integer T_s and $c \in (0,1)$ are some constants.

	iteration complexity	rate
SSGD [Avron *et al.*, 2012]	$O(mk_t^2)$	$O(1/\sqrt{T})$
active subspace selection [Hsieh and Olsen, 2014]	$O(\|\Omega\|_1 k_t^2 T_a)$	$O(c^{T-T_s})$
LMaFit [Wen *et al.*, 2012]	$O(\|\Omega\|_1 k_t + mk_t)$	—
boosting [Zhang *et al.*, 2012]	$O(\|\Omega\|_1 t^2 T_a)$	$O(1/T)$
TR [Mishra *et al.*, 2013]	$O(\|\Omega\|_1 t^2 T_a)$	—
ALT-Impute [Hastie *et al.*, 2015]	$O(\|\Omega\|_1 k_t + mk_t^2)$	$O(1/T)$
AIS-Impute	$O((r_t+k_t)\|\Omega\|_1 + (r_{t-1}+r_t+k_t)mk_t)$	$O(1/T^2)$

square loss is considered in [Mazumder *et al.*, 2010]. Here, we extend it to any smooth convex loss.

Let the rank-k matrix obtained from Algorithm 5 be $X = U\Sigma V^\top$, where $U = [u_i]$ and $V = [v_i]$. Let $A(\theta) = U\mathrm{Diag}(\theta)V^\top$. We undo part of the shrinkage on the singular values by replacing X with $A(\theta_*)$, where

$$\theta_* = \arg\min_\theta \phi(\theta) \equiv \sum_{(i,j)\in\Omega} \ell(A(\theta)_{ij}, O_{ij}). \tag{22}$$

When ℓ is the square loss, (22) has a closed-form solution [Mazumder *et al.*, 2010]. However, for smooth convex ℓ in general, this is not possible and we optimize (22) using L-BFGS [Nocedal and Wright, 1999]. The most expensive step in each L-BFGS iteration is the computation of the gradient $\nabla\phi(\theta) \in \mathbb{R}^k$, where $[\nabla\phi(\theta)]_i = u_i^\top B v_i$, $B_{ij} = \frac{d\ell(A(\theta)_{ij}, O_{ij})}{dA(\theta)_{ij}}$ if $(i,j) \in \Omega$, and 0 otherwise. As S is sparse, computing $\nabla\phi(\theta)$ only takes $O(k\|\Omega\|_1)$ time. Moreover, L-BFGS has superlinear convergence [Nocedal and Wright, 1999]. Empirically, it converges in fewer than 10 iterations.

4 Tensor Completion

Complicated data objects can often be arranged as tensors. In this section, we extend the proposed Algorithm 5 in Section 3 from matrices to tensors.

4.1 Tensor Model

The nuclear norm can be defined on tensors in various ways. The following two are the most popular.

Definition 4.1 ([Tomioka *et al.*, 2010])**.** *For a D-order tensor $\mathcal{X}$, the* overlapped nuclear norm *is*

$$\|\mathcal{X}\|_{\textit{overlap}} = \sum_{d=1}^D \lambda_d \|\mathcal{X}_{\langle d\rangle}\|_*,$$

and the scaled latent nuclear norm *is*

$$\|\mathcal{X}\|_{\textit{scaled}} = \min_{\mathcal{X}^1,\ldots,\mathcal{X}^D \,:\, \sum_{d=1}^D \mathcal{X}^d = \mathcal{X}} \sum_{d=1}^D \lambda_d \|\mathcal{X}^d_{\langle d\rangle}\|_*.$$

Here, $\lambda_d \geq 0$'s are hyperparameters.

The overlapped nuclear norm regularizer penalizes nuclear norms on all modes. When only several modes are low-rank, decomposition with the scaled latent nuclear norm has better generalization [Tomioka *et al.*, 2010; Wimalawarne *et al.*, 2014]. In this paper, we focus on the scaled latent nuclear norm regularizer.

Given a partially observed tensor $\mathcal{O} \in \mathbb{R}^{I_1\times\cdots\times I_D}$, with the observed entries indicated by $\Omega \in \{0,1\}^{I_1\times\cdots\times I_D}$. The tensor completion problem can be formulated as

$$\min_{\mathcal{X}^1,\ldots,\mathcal{X}^D} F\left([\mathcal{X}_t^1,\ldots,\mathcal{X}_t^D]\right) \tag{23}$$
$$\equiv \sum_{(i_1,\ldots,i_D)\in\Omega} \ell\left(\sum_{d=1}^D \mathcal{X}^d_{i_1\ldots i_D}, \mathcal{O}_{i_1\ldots i_D}\right) + \sum_{d=1}^D \lambda_d \|\mathcal{X}^d_{\langle d\rangle}\|_*.$$

The recovered tensor is $\mathcal{X} = \sum_{d=1}^D \mathcal{X}^d$. In [Tomioka *et al.*, 2010; Liu *et al.*, 2013], problem (23) is solved using ADMM [Boyd *et al.*, 2011]. However, the ADMM update involves SVD in each iteration, which takes $O\left(\prod_{d=1}^D I_d \sum_{d=1}^D I_d\right)$ time and can be expensive.

4.2 Generalizing SVT

In (23), let

$$f([\mathcal{X}^1,\ldots,\mathcal{X}^D]) = \sum_{(i_1,\ldots,i_D)\in\Omega} \ell\left(\sum_{d=1}^D \mathcal{X}^d_{i_1\ldots i_D}, \mathcal{O}_{i_1\ldots i_D}\right), \tag{24}$$

$$g([\mathcal{X}^1,\ldots,\mathcal{X}^D]) = \sum_{d=1}^D \lambda_d \|\mathcal{X}^d_{\langle d\rangle}\|_*. \tag{25}$$

The iterates in Algorithm 5 are generated by SVT. As there are multiple nuclear norms in (25), the following extends SVT for this case.

As g in (25) is separable w.r.t. $\mathcal{X}^i$'s, one can compute the proximal step for each $\mathcal{X}^i$ separately [Parikh and Boyd, 2014]. Updates (5), (6) in the accelerated proximal gradient algorithm (Algorithm 1) become

$$\begin{aligned} \mathcal{Y}_t^d &= (1+\theta_t)\mathcal{X}_t^d - \theta_t \mathcal{X}_{t-1}^d, \\ \breve{\mathcal{Z}}_t^d &= \mathcal{Y}_t^d - \mu\mathcal{S}_t = (1+\theta_t)\mathcal{X}_t^d - \theta_t\mathcal{X}_{t-1}^d - \mu\mathcal{S}_t, \end{aligned} \tag{26}$$

for $d = 1,\ldots,D$, where $\mathcal{S}_t$ is a sparse tensor with

$$(\mathcal{S}_t)_{i_1\ldots i_D} = \begin{cases} \frac{d\ell((\hat{\mathcal{Y}}_t)_{i_1\ldots i_D}, \mathcal{O}_{i_1\ldots i_D})}{d(\hat{\mathcal{Y}}_t)_{i_1\ldots i_D}} & \text{if } (i_1,\ldots,i_D)\in\Omega \\ 0 & \text{otherwise} \end{cases}, \tag{27}$$

and $\hat{\mathcal{Y}}_t = \sum_{d=1}^{D} \mathcal{Y}_t^d$. Lemma 2.2 is also extended to $[\mathcal{X}_{t+1}^1, \ldots, \mathcal{X}_{t+1}^D] = \text{prox}_{\mu g}([\breve{\mathcal{Z}}_t^1, \ldots, \breve{\mathcal{Z}}_t^D])$ as follows. The proof can be found in Appendix A.5.

Proposition 4.1. $(\mathcal{X}_{t+1}^d)_{\langle d \rangle} = SVT_{\mu\lambda_d \|\cdot\|_*}\left((\breve{\mathcal{Z}}_t^d)_{\langle d \rangle}\right)$.

The stepsize rule in (4) depends on the modulus of Lipschitz smoothness of f, which is given by the following. The proof can be found in Appendix A.6.

Proposition 4.2. *If ℓ is ρ-Lipschitz smooth, f in* (24) *is $\sqrt{D}\rho$-Lipschitz smooth.*

Proposition 3.3 can be used to reduce the size of $(\breve{\mathcal{Z}}_t^d)_{\langle d \rangle}$ in Proposition 4.1, and Algorithm 3 can be used to approximate the underlying SVD. However, this is still not fast enough. Assume that $\breve{k}_t^d$ singular values in $(\breve{\mathcal{Z}}_t^d)_{\langle d \rangle}$ are larger than $\mu\lambda_d$, and rank-k_t^d SVD, where $k_t^d \geq \breve{k}_t^d$, is performed. SVT on $(\mathcal{Z}_t^d)_{\langle d \rangle}$ takes $O\left(k_t^d \prod_{d=1}^{D} I_d\right)$ time. As SVT has to be performed on each mode, one iteration of Algorithm 1 takes $O\left(\prod_{d=1}^{D} I_d \sum_{d=1}^{D} k_t^d\right)$ time, which is expensive.

4.3 Fast Approximate SVT with Special Structure

In Section 3.2, the special "sparse plus low-rank" structure can greatly reduce the time complexity of matrix multiplications. As $\mathcal{X}_{t-1}^d, \mathcal{X}_t^d$ are low-rank tensors and $\mathcal{S}_t$ is sparse, $\breve{\mathcal{Z}}_t^d$ in (26) also has the "sparse plus low-rank" structure. However, to generate $\mathcal{X}_{t+1}^d$ using Proposition 4.1, we need to perform matrix multiplications of the form $(\breve{\mathcal{Z}}_t^d)_{\langle d \rangle} v$, where $v \in \mathbb{R}^{I_{D \setminus d}}$, and $u^\top (\breve{\mathcal{Z}}_t)_{\langle d \rangle}$, where $u \in \mathbb{R}^{I_d}$. Unfolding $\breve{\mathcal{Z}}_t$ takes $O\left(\prod_{d=1}^{D} I_d\right)$ time and can be expensive. In the following, we show how this can be avoided.

To generate $(\mathcal{X}_{t+1}^d)_{\langle d \rangle}$, it can be seen from Proposition 4.1 and (26) that $\mathcal{X}_t^d$ and $\mathcal{X}_{t-1}^d$ only need to be unfolded along their dth modes. Hence, instead of storing them as tensors, we store $(\mathcal{X}_t^d)_{\langle d \rangle}$ as its rank-r_t^d SVD $U_t^d \Sigma_t^d {V_t^d}^\top$, and $(\mathcal{X}_{t-1}^d)_{\langle d \rangle}$ as its rank-r_{t-1}^d SVD $U_{t-1}^d \Sigma_{t-1}^d {V_{t-1}^d}^\top$. For any $v \in \mathbb{R}^{I_{D \setminus d}}$,

$$\begin{aligned}(\breve{\mathcal{Z}}_t^d)_{\langle d \rangle} v \;=\; & (1+\theta_t) U_t^d \Sigma_t^d ({V_t^d}^\top v) \\ & -\theta_t U_{t-1}^d \Sigma_{t-1}^d ({V_{t-1}^d}^\top v) - \mu(\mathcal{S}_t)_{\langle d \rangle} v.\end{aligned}$$

The first two terms can be computed in $O((I_d + I_{D \setminus d})(r_t^d + r_{t-1}^d))$ time. As $\mathcal{S}_t$ is sparse, computing the last term takes $O(\|\Omega\|_1)$ time. Thus, $(\breve{\mathcal{Z}}_t^d)_{\langle d \rangle} v$ can be obtained in $O(\|\Omega\|_1 + (I_d + I_{D \setminus d})(r_t^d + r_{t-1}^d))$ time. Similarly, for any $u \in \mathbb{R}^{I_d}$, $u^\top (\breve{\mathcal{Z}}_t)_{\langle d \rangle}$ can be computed in $O(\|\Omega\|_1 + (I_d + I_{D \setminus d})(r_t^d + r_{t-1}^d))$ time. Thus, performing approximate SVT on $(\breve{\mathcal{Z}}_t^d)_{\langle d \rangle}$, with rank $k_t^d \geq \breve{k}_t^d$, using Algorithm 4 takes $O(k_t^d \|\Omega\|_1 + k_t^d (I_d + I_{D \setminus d})(r_t^d + r_{t-1}^d))$ time. Using Proposition 4.1, solving the proximal step $\text{prox}_{\mu g}([\breve{\mathcal{Z}}_t^1, \ldots, \breve{\mathcal{Z}}_t^D])$ takes a total of

$$O\left(\sum_{d=1}^{D} k_t^d \|\Omega\|_1 + k_t^d (I_d + I_{D \setminus d})(r_t^d + r_{t-1}^d)\right) \tag{28}$$

time. As the target tensor is low-rank, $r_t^d, k_t^d \ll I_d$ for $d = 1, \ldots, D$. Hence, (28) is much faster than directly using Proposition 4.1 ($O(\prod_{d=1}^{D} I_d \sum_{d=1}^{D} k_t^d)$ time).

4.4 The Proposed Algorithm

The whole procedure is shown in Algorithm 6. Unlike, Algorithm 5, D SVTs have to be computed (steps 5-11) in each iteration.

Algorithm 6 AIS-Impute (tensor case).

Require: partially observed tensor $\mathcal{O}$, parameter λ;
1: initialize $c = 1$, $\mathcal{X}_0^1 = \cdots = \mathcal{X}_0^D = 0$, $\mathcal{X}_1^1 = \cdots = \mathcal{X}_1^D = 0$, step-size $\mu = 1/(\sqrt{D}\rho)$, $\hat{\lambda} > \max_{d=1,\ldots,D} \lambda_d$ and $\nu \in (0,1)$;
2: **for** $t = 1, 2, \ldots, T$ **do**
3: $\quad \theta_t = (c-1)/(c+2)$;
4: $\quad$ construct the sparse observed tensor $\mathcal{S}_t$ from (27);
5: $\quad$ **for** $d = 1, \ldots, D$ **do**
6: $\quad\quad (\lambda_d)_t = (\hat{\lambda} - \lambda_d)\nu^{t-1} + \lambda_d$;
7: $\quad\quad \breve{\mathcal{Z}}_t^d = (1+\theta_t)\mathcal{X}_t^d - \theta_t \mathcal{X}_{t-1}^d + \mu \mathcal{S}_t$;
8: $\quad\quad V_{t-1}^d = V_{t-1}^d - V_t^d \left((V_t^d)^\top V_{t-1}^d\right)$, remove zero columns;
9: $\quad\quad R_t^d = \text{QR}(\left[V_t^d, V_{t-1}^d\right])$;
10: $\quad\quad \left[U_{t+1}^d, \Sigma_{t+1}^d, V_{t+1}^d\right] =$ approx-SVT $\left((\breve{\mathcal{Z}}_t^d)_{\langle d \rangle}, R_t^d, \mu(\lambda_d)_t, J\right)$; // $\mathcal{X}_{\langle d \rangle}^d = U_{t+1}^d \Sigma_{t+1}^d (V_{t+1}^d)^\top$
11: $\quad$ **end for**
12: $\quad$ **if** $F\left([\mathcal{X}_{t+1}^1, \ldots, \mathcal{X}_{t+1}^D]\right) > F\left([\mathcal{X}_t^1, \ldots, \mathcal{X}_t^D]\right)$ **then**
13: $\quad\quad c = 1$;
14: $\quad$ **else**
15: $\quad\quad c = c + 1$;
16: $\quad$ **end if**
17: **end for**
18: **return** $U_{t+1}^d, \Sigma_{t+1}^d, V_{t+1}^d$ where $d = 1, \ldots, D$.

Analogous to Theorem 3.5, we have the following. The proof can be found in Appendix A.7.

Theorem 4.3. *Assume that (i) $k_t^d \geq \breve{k}_t^d$ for $d = 1, \ldots, D$, all t and $J = t$; (ii) $F\left([\mathcal{X}_t^1, \ldots, \mathcal{X}_t^D]\right)$ is upper bounded. The sequence $\left\{[\mathcal{X}_t^1, \ldots, \mathcal{X}_t^D]\right\}$ generated from Algorithm 6 converges to the optimal solution with a $O(1/T^2)$ rate.*

4.5 Post-Processing

As in Section 3.6, the nuclear norm regularizer in (23) may over-penalize top singular values. To undo such shrinkage and boost recovery performance, we also adopt post-processing here. Let the tensor output from Algorithm 6 be $\mathcal{X} = \sum_{d=1}^{D} \mathcal{X}^d$, where $\mathcal{X}_{\langle d \rangle}^d = U^d \Sigma^d (V^d)^\top$ has rank k^d.

Define $\mathcal{A}\left(\theta^1,\ldots,\theta^D\right) = \sum_{d=1}^{D}\left(U^d\text{Diag}(\theta^d)(V^d)^\top\right)_{\langle d\rangle}$. As in (22), we replace $\mathcal{X}$ with $\mathcal{A}\left(\theta_*^1,\ldots,\theta_*^D\right)$, where

$$\left[(\theta_*^1)^\top,\ldots,(\theta_*^D)^\top\right]^\top = \arg\min_{\theta^1,\ldots,\theta^D}\phi(\theta^1,\ldots,\theta^D), \quad (29)$$

and

$$\phi(\theta^1,\ldots,\theta^D) = \sum_{(i_1,\ldots,i_D)\in\Omega}\ell(\mathcal{A}(\theta^1,\ldots,\theta^D)_{i_1\ldots i_D},\mathfrak{O}_{i_1\ldots i_D}).$$

As (29) is a smooth convex problem, L-BFGS is used for optimization. Let $U^d = [u^d]$ and $V^d = [v^d]$. Then, $\nabla\phi(\theta^1,\ldots,\theta^D) = \left[(w^1)^\top,\ldots,(w^D)^\top\right]^\top$ where $w^d = [w_i^d] \in k^d$, $w_i^d = (u_i^d)^\top\mathcal{B}_{\langle d\rangle}v_i^d$, and $\mathcal{B}_{i_1\ldots i_D} = \frac{d\ell\left(\mathcal{A}([\theta^1,\ldots,\theta^D])_{i_1\ldots i_D},\mathfrak{O}_{i_1\ldots i_D}\right)}{d\mathcal{A}([\theta^1,\ldots,\theta^D])_{i_1\ldots i_D}}$ if $(i_1,\ldots,i_D)\in\Omega$ and 0 otherwise. Computation of $\nabla\phi(\theta^1,\ldots,\theta^D)$ takes $O(\sum_{d=1}^{D}k^d\|\Omega\|_1)$ time, and each L-BFGS iteration is cheap.

5 Experiments

In this section, we perform experiments on matrix completion (Sections 5.1-5.3) and tensor completion (Sections 5.4, 5.5). Experiments are performed on a PC with Intel Xeon E5-2695 CPU and 256GB RAM. All algorithms are implemented in Matlab, with operations on Ω written in C.

5.1 Synthetic Data

In this section, we perform matrix completion experiments using synthetic data. The ground-truth matrix has a rank of 5, and is generated as $O = UV \in \mathbb{R}^{m\times m}$, where the entries of $U \in \mathbb{R}^{m\times 5}$ and $V \in \mathbb{R}^{5\times m}$ are sampled i.i.d. from the standard normal distribution $\mathcal{N}(0,1)$. Noise, sampled from $\mathcal{N}(0,0.05)$, is then added. We randomly choose $15m\log(m)$ of the entries in O as observed. Half of them are used for training, and the other half as validation set for parameter tuning. Testing is performed on the unobserved (missing) entries. We vary m in the range $\{500,1000,2000\}$.

The following proximal algorithms are compared:

1. accelerated proximal gradient algorithm[2] (denoted "APG") [Toh and Yun, 2010]: It uses PROPACK to obtain singular values that are larger than λ;
2. Soft-Impute[3] [Mazumder *et al.*, 2010];
3. AIS-Impute (the proposed Algorithm 5);
4. AIS-Impute (exact): This is a variant of the proposed algorithm with exact SVT step (computed using PROPACK).

Let X be the recovered matrix. For performance evaluation, we use the (i) normalized mean squared error NMSE $= \|P_{\Omega^\perp}(X-UV)\|_F/\|P_{\Omega^\perp}(UV)\|_F$, and (ii) rank of X. To reduce statistical variability, experimental results are averaged over 5 repetitions.

Results are shown in Table 3. As can be seen, all algorithms have similar NMSE performance, with Soft-Impute being slightly worse. The comparison of objective vs time and iterations are shown in Figure 1. In terms of the number of iterations, the accelerated algorithms (APG, AIS-Impute(exact) and AIS-Impute) are very similar and converge much faster than Soft-Impute (which only has a $O(1/T)$ convergence rate). However, in terms of time, both APG and Soft-Impute are slow, as APG does not utilize the "sparse plus low-rank" structure and Soft-Impute has slow convergence. AIS-Impute(exact) is consistently faster than APG and Soft-Impute, as both acceleration and "sparse plus low-rank" structure are utilized. However, AIS-Impute is the fastest as it further allows inexact updates of the proximal step. This also verifies our motivation of using approximate SVT at Section 3.3.

Table 3: Matrix completion results on synthetic data (here, sparsity is the proportion of observed entries). No post-processing is performed and NMSE is scaled by 10^{-3}. The lowest and comparable NMSEs (according to the pairwise t-test with 95% confidence) are highlighted.

		NMSE	rank
$m = 500$ sparsity: 18.6%	APG	17.1±0.1	5
	Soft-Impute	17.0±0.1	5
	AIS-Impute (exact)	**16.8±0.1**	5
	AIS-Impute	**16.9±0.1**	5
$m = 1000$ sparsity: 10.4%	APG	**16.7±0.1**	5
	Soft-Impute	17.5±0.1	5
	AIS-Impute (exact)	16.8±0.1	5
	AIS-Impute	**16.7±0.1**	5
$m = 2000$ sparsity: 5.7%	APG	**14.3±0.1**	5
	Soft-Impute	14.9±0.1	5
	AIS-Impute (exact)	**14.2±0.1**	5
	AIS-Impute	**14.2±0.1**	5

Table 4 shows the NMSE results with post-processing in Section 3.6. Compared with the time used by the main algorithm (Figure 1), the post-processing time is small and can be ignored. Thus, post-processing are always be performed in Sections 5.2 and 5.3.

5.2 Recommender System

In this section, we consider the standard matrix completion problem in (1). Experiments are performed on two well-known benchmark data sets, *MovieLens* (Section 5.2) and *Netflix* (Section 5.2).

MovieLens

The *MovieLens* data set[4] (Table 5) contains ratings (in $\{1,2,3,4,5\}$) of different users on movies, It has been commonly used in matrix completion experiments [Mazumder *et al.*, 2010; Hsieh and Olsen, 2014]. We randomly use 50% of the observed ratings for training, 25% for validation and the rest for testing.

Besides the proximal gradient algorithms in Section 5.1, we also compare with the following state-of-the-art non-proximal matrix completion algorithms:

[2] http://www.math.nus.edu.sg/~mattohkc/NNLS.html

[3] http://www.mit.edu/~rahulmaz/software.html

[4] http://grouplens.org/datasets/movielens/

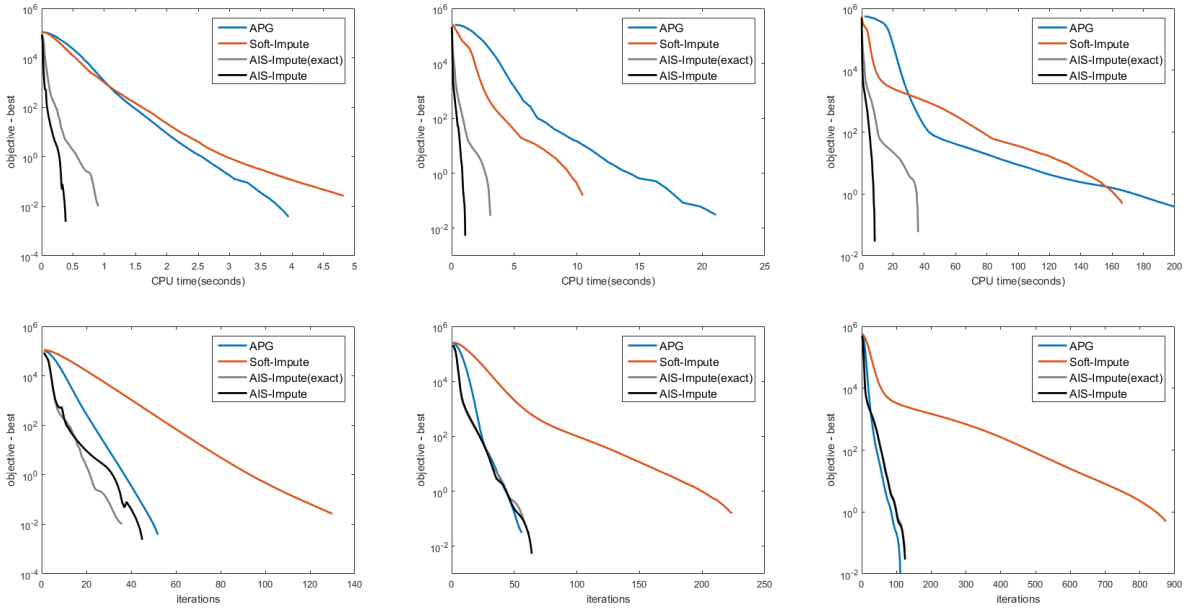

(a) $m = 500$. (b) $m = 1000$. (c) $m = 2000$.

Figure 1: Convergence of objective value on the synthetic matrix data. Top: vs CPU time (in seconds); Bottom: vs number of iterations.

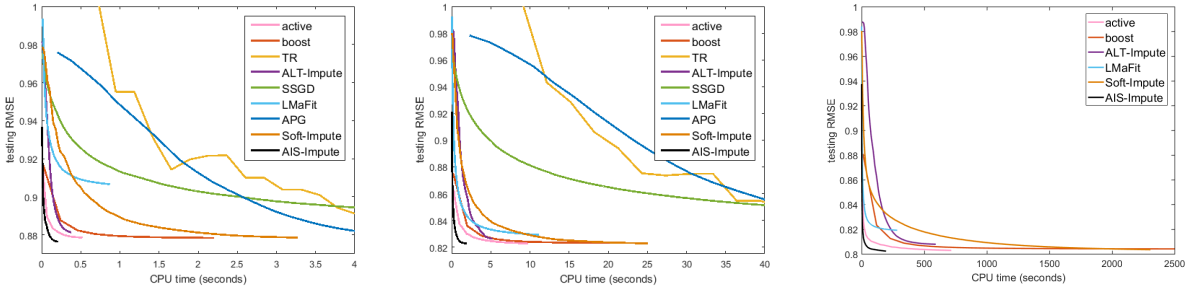

(a) *100K*. (b) *1M*. (c) *10M*.

Figure 2: Testing RMSE vs CPU time (in seconds) on the *MovieLens* data sets.

1. active subspace selection (denoted "active")[5] algorithm [Hsieh and Olsen, 2014]: In each iteration, this algorithm uses the power method to identify the active row and column subspaces, and then reduces the nuclear norm optimization problem to a smaller problem;
2. boosting [Zhang *et al.*, 2012][6] , a variant of the Frank-Wolfe algorithm [Frank and Wolfe, 1956] for matrix completion. To speedup convergence local optimization (using L-BFGS) is performed in each iteration;
3. second-order trust-region algorithm (denoted "TR")[7], which alternates between fixed-rank optimization and rank-one updates [Mishra *et al.*, 2013];
4. ALT-Impute [Hastie *et al.*, 2015], a speedup variant of Soft-Impute that avoids SVD by alternating least squares. Its convergence rate is $O(1/T)$;
5. SSGD [Avron *et al.*, 2012], a stochastic algorithm for the nuclear norm regularized matrix completion problem; and
6. a fixed-rank approach solving by alternative minimization (denoted "LMaFit") [8]. Over-relaxation is used for further speedup [Wen *et al.*, 2012].

Except LMaFit, all the other algorithms solve (1) or its equivalent problem. Let X be the recovered matrix,

[5] http://www.cs.utexas.edu/~cjhsieh/nuclear_active_1.1.zip

[6] http://users.cecs.anu.edu.au/~xzhang/GCG/

[7] https://sites.google.com/site/bamdevm/codes/tracenorm

[8] http://www.caam.rice.edu/~optimization/L1/LMaFit/download.html

Table 4: Matrix completion results on synthetic data, with post-processing. NMSE is scaled by 10^{-3}. The lowest and comparable NMSEs (according to the pairwise t-test with 95% confidence) are highlighted.

		NMSE	post-proc time (sec)
$m = 500$	APG	**9.8±0.1**	0.01
	Soft-Impute	9.9±0.1	0.01
	AIS-Impute (exact)	**9.8±0.1**	0.01
	AIS-Impute	**9.8±0.1**	0.01
$m = 1000$	APG	**9.1±0.1**	0.01
	Soft-Impute	9.9±0.1	0.03
	AIS-Impute (exact)	9.4±0.1	0.02
	AIS-Impute	9.3±0.1	0.02
$m = 2000$	APG	8.5±0.1	0.05
	Soft-Impute	9.3±0.1	0.18
	AIS-Impute (exact)	**8.4±0.1**	0.04
	AIS-Impute	**8.4±0.1**	0.04

Table 5: *MovieLens* data sets used in the experiments.

	#users	#movies	# observed ratings
100K	943	1,682	100,000
1M	6,040	3,449	999,714
10M	69,878	10,677	10,000,054

and the testing ratings $\{\hat{O}_{ij}\}$ be indexed by the set $\hat{\Omega}$. For performance evaluation, as in [Hsieh and Olsen, 2014; Mazumder *et al.*, 2010], we use (i) the testing root mean squared error RMSE $= \sqrt{\|P_{\hat{\Omega}}(X - \hat{O})\|_F^2 / \|\hat{\Omega}\|_1}$; and (ii) rank of X. The experiment is repeated 5 times and the average performance is reported.

Results are shown in Table 6. As can be seen, AIS-Impute is consistently the fastest and has the lowest RMSE. On *MovieLens-10M*, TR, SSGD and APG are not run as they are too slow. Figure 2 shows the testing RMSE with CPU time. As can be seen, Boost, TR, SSGD and APG are all very slow. Boost and TR need to solve an expensive subproblem in each iteration; SSGD has slow convergence; while APG requires SVD and does not utilize the "sparse plus low-rank" structure for fast matrix multiplication. ALT-Impute and LMaFit do not need SVT, and are faster than Soft-Impute. However, their nonconvex formulations have slow convergence, and are thus slower than AIS-Impute. Overall, AIS-Impute is the fastest, as it combines cheap iteration and fast convergence.

Netflix

In this Section, we demonstrate the speedup of AIS-Impute over other algorithms solving the nuclear norm regularized problem (1) on the *Netflix* data set. It contains ratings of 480,189 users on 17,770 movies. 1% of the ratings matrix are observed. We randomly sample 50% of the observed ratings for training, and the rest for testing.

We only compare with active subspace selection, ALT-Impute and Soft-Impute; while methods including boosting, TR, SSGD, APG are slow and not compared. LMaFit solves an a different optimization problem based on matrix factorization, and has worse recovery performance than AIS-Impute. Thus, it is also not compared. As in [Mazumder *et al.*, 2010], several choices of λ are experimented.

Table 6: Results on the *MovieLens* data sets. Algorithms that cannot converge in 10^4 seconds are not reported. The lowest and comparable RMSEs (according to the pairwise t-test with 95% confidence) are highlighted.

		RMSE	rank
100K	active	**0.880±0.003**	8
	boost	**0.881±0.003**	8
	TR	0.884±0.002	8
	ALT-Impute	**0.882±0.003**	8
	SSGD	0.886±0.011	8
	LMaFit	0.896±0.011	3
	APG	**0.880±0.003**	8
	Soft-Impute	**0.881±0.003**	8
	AIS-Impute	**0.880±0.003**	8
1M	active	**0.821±0.001**	16
	boost	**0.821±0.001**	16
	TR	**0.821±0.001**	16
	ALT-Impute	0.824±0.002	16
	SSGD	0.841±0.009	16
	LMaFit	0.827±0.001	6
	APG	**0.820±0.001**	16
	Soft-Impute	**0.821±0.001**	16
	AIS-Impute	**0.820±0.002**	16
10M	active	**0.803±0.001**	74
	boost	0.804±0.001	74
	ALT-Impute	0.808±0.001	74
	LMaFit	0.818±0.001	12
	Soft-Impute	0.804±0.001	74
	AIS-Impute	**0.802±0.001**	74

Results are shown in Table 7. As in previous experiments, the RMSEs and ranks obtained by the various algorithms are similar. Figure 3 shows the detailed comparison on testing RMSE versus CPU time. As can be seen, AIS-Impute is again much faster.

5.3 Link Prediction

Given a graph with m nodes and an incomplete adjacency matrix $O \in \{\pm 1\}^{m \times m}$, link prediction aims to recover a low-rank matrix $X \in \mathbb{R}^{m \times m}$ such that the signs of X_{ij}'s and O_{ij}'s agree on most of the observed entries. This is a binary matrix completion problem [Chiang *et al.*, 2014], and we use the logistic loss $\ell(X_{ij}, O_{ij}) \equiv \log(1 + \exp(-X_{ij}O_{ij}))$ in (19).

Experiments are performed on the *Epinions* and *Slashdot* data sets[9] [Chiang *et al.*, 2014] (Table 8). Each row/column of the matrix O corresponds to a user (users with fewer than two observations are removed). For *Epinions*, $O_{ij} = 1$ if user i trusts user j, and -1 otherwise. Similarly for *Slashdot*, $O_{ij} = 1$ if user i tags user j as friend, and -1 otherwise. As can be seen from previous sections, Boost, TR, SSGD, APG and Soft-Impute are all slow, and thus

[9] https://snap.stanford.edu/data/

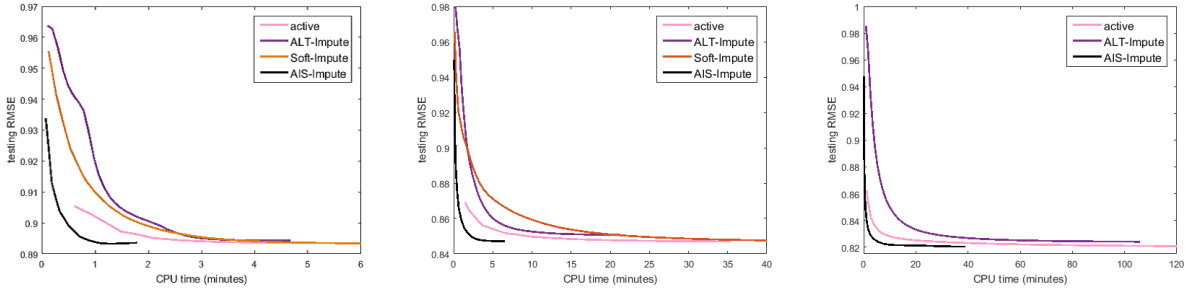

(a) $\lambda = \lambda_0/10$. (b) $\lambda = \lambda_0/20$. (c) $\lambda = \lambda_0/30$.

Figure 3: Testing RMSE vs CPU time (in minutes) on the *Netflix* data set, with various values for the regularization parameter λ.

Table 7: Results on the *Netflix* data set. The regularization parameter λ in (1) is set as λ_0/c, where $\lambda_0 = \|P_\Omega(O)\|_F$. The lowest and comparable RMSEs (according to the pairwise t-test with 95% confidence) are highlighted. Soft-Impute with $c = 30$ is not run as it is very slow.

		RMSE	rank
$c = 10$	active	0.894±0.001	3
	ALT-Impute	0.900±0.006	3
	Soft-Impute	**0.893±0.001**	3
	AIS-Impute	**0.893±0.001**	3
$c = 20$	active	**0.847±0.001**	14
	ALT-Impute	0.850±0.001	14
	Soft-Impute	**0.847±0.001**	14
	AIS-Impute	**0.847±0.001**	14
$c = 30$	active	**0.820±0.001**	116
	ALT-Impute	0.825±0.001	116
	AIS-Impute	**0.820±0.001**	116

they are not considered here. Besides, LMaFit and ALT-Impute are designed for the square loss. Thus, comparison is performed with (i) active subspace selection; (ii) AIS-Impute; and (iii) AltMin: the alternative minimization approach used in [Chiang *et al.*, 2014]. We use 80% of the ratings for training, 10% for validation and the rest for testing. Let X be the recovered matrix, and the test set $\{\hat{O}_{ij}\}$ be indexed by the set $\hat{\Omega}$. For performance evaluation, we use the (i) testing accuracy $\frac{1}{\|\hat{\Omega}\|_1}\sum_{(i,j)\in\hat{\Omega}} I(\text{sign}(X_{ij}) = \hat{O}_{ij})$, where $I(\cdot)$ is the indicator function; and (ii) the rank of X. To reduce statistical variability, experimental results are averaged over 5 repetitions.

Table 8: Datasets used for link prediction.

	#rows	#columns	#signs
Epinions	84,601	48,091	505,074
Slashdot	70,284	32,188	324,745

Results are shown in Table 9 and Figure 4 shows the testing accuracy with CPU time. As can be seen, active and AIS-Impute have slightly better accuracies than AltMin, and AIS-Impute is the fastest.

Table 9: Performance for link prediction in social network. The lowest and comparable RMSEs (according to the pairwise t-test with 95% confidence) are highlighted.

		accuracy	rank
Epinions	active	**0.939±0.002**	12
	AltMin	0.936±0.002	41
	AIS-Impute	**0.940±0.001**	12
Slashdot	active	**0.844±0.001**	16
	AltMin	0.839±0.002	39
	AIS-Impute	**0.843±0.001**	16

5.4 Tensor Completion: Synthetic Data

In this section, we perform tensor completion experiments using synthetic data. The ground-truth data tensor (of size $m \times m \times 3$) is generated as $\mathcal{O} = \mathcal{C} \times_1 A_1 \times_2 A_2 \times_3 A_3$, where the elements of $A_1 \in \mathbb{R}^{m\times 3}$, $A_2 \in \mathbb{R}^{m\times 3}$, $A_3 \in \mathbb{R}^{3\times 3}$ and the core tensor $\mathcal{C} \in \mathbb{R}^{3\times 3}$ are all sampled i.i.d. from the standard normal distribution $\mathcal{N}(0, 1)$, and $\times_k$ is the k-mode product [10]. Thus, $\mathcal{O}$ is low-rank on the first two mode but not on the third one. Noise $\mathcal{G}$, with its elements sampled i.i.d. from the normal distribution $\mathcal{N}(0, 0.05)$, is then added. A total number of $\Omega = 45m\log(m)$ random elements in $\mathcal{O}$ are observed. Half of them are used for training, and the other half for validation. On testing, we perform evaluation on the unobserved entries and use the same criteria as in Section 5.1, i.e., NMSE and recovered rank on each mode.

Similar to Section 5.1, we compare the following algorithms: (i) APG; (ii) the proposed algorithm with exact SVD (AIS-Impute(exact)), and (iii) the proposed algorithm which uses power method to approximate SVT (AIS-Impute). Soft-Impute has not been extended to tensor completion, and is thus not compared. Besides, we also compare with the ADMM approach in [Tomioka *et al.*, 2010] (denoted "ADMM(scaled)").

[10] The k-mode product of a tensor $\mathcal{X}$ and a matrix A is defined as $\mathcal{X} \times_k A = \left(\mathcal{X}_{\langle k\rangle} A\right)_{\langle k\rangle}$ [Kolda and Bader, 2009].

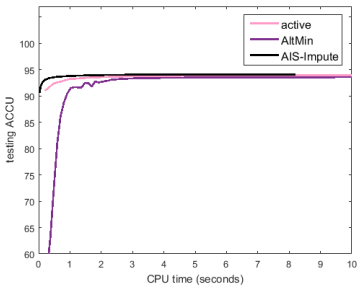

(a) *Epinions*.

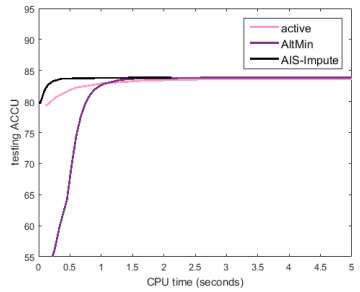

(b) *Slashdot*.

Figure 4: Testing accuracy vs CPU time (in seconds) on the *Epinions* and *Slashdot* data sets.

As suggested in [Wimalawarne *et al.*, 2014], we set $(\lambda_1, \lambda_2, \lambda_3)$ in the scaled latent nuclear norm to $(1, 1, \frac{\sqrt{m}}{\sqrt{3}})\lambda$. Thus, the only tunable parameter is λ, which is obtained by grid search using the validation set. We also vary m in $\{500, 1000, 2000\}$. Experimental results are averaged over 5 repetitions.

Results on NMSE and rank are shown in Table 10. As can be seen, APG, AIS-Impute(exact) and AIS-Impute have comparable performance. The comparison of objective vs time and iterations are shown in Figure 5. In terms of iterations, APG, AIS-Impute(exact) and AIS-Impute have similar behavior as they all have $O(1/T^2)$ convergence rate. These also agree with the matrix case in Section 5.1. They are faster than ADMM(scaled), which only has a slower $O(1/T)$ rate [He and Yuan, 2012]. In terms of time, as APG does not utilize the "sparse plus low-rank" structure, it is slower than AIS-Impute(exact) and AIS-Impute. ADMM(scale) is even slower than APG, as it only has slower convergence rate compared with APG. AIS-Impute is the fastest, as it has both fast $O(1/T^2)$ convergence rate and low per-iteration complexity.

Performance with post-processing in Section 4.5 is shown in Table 11. As can be seen, it is very efficient and improves NMSE. Thus, we always perform post-processing in Section 5.5.

Table 10: Tensor completion results on synthetic data. No post-processing is performed and NMSE is scaled by 10^{-3}. The lowest and comparable NMSEs (according to the pairwise t-test with 95% confidence) are highlighted.

		NMSE	rank of mode		
			1	2	3
$m = 500$ sparsity: 18.6%	APG	**17.3±1.4**	3	3	0
	AIS-Impute(exact)	**17.3±1.8**	3	3	0
	AIS-Impute	**17.2±1.7**	3	3	0
	ADMM(scaled)	**17.5±1.8**	3	3	0
$m = 1000$ sparsity: 10.4%	APG	**17.2±1.3**	3	3	0
	AIS-Impute(exact)	**17.1±1.5**	3	3	0
	AIS-Impute	**17.0±1.4**	3	3	0
	ADMM(scaled)	**17.0±2.3**	3	3	0
$m = 2000$ sparsity: 5.7%	APG	**18.1±2.3**	3	3	0
	AIS-Impute(exact)	**18.0±1.8**	3	3	0
	AIS-Impute	**18.0±1.7**	3	3	0
	ADMM(scaled)	**18.3±2.3**	3	3	0

Table 11: Tensor completion results on synthetic data (with post-processing). NMSE is scaled by $\times 10^{-3}$. The lowest and comparable NMSEs (according to the pairwise t-test with 95% confidence) are highlighted.

		NMSE	post-proc time (sec)
$m = 500$	APG	**10.5±0.4**	0.1
	AIS-Impute(exact)	**10.5±0.4**	0.1
	AIS-Impute	**10.5±0.5**	0.1
	ADMM(scaled)	**10.5±0.6**	0.1
$m = 1000$	APG	**11.2±0.4**	0.1
	AIS-Impute(exact)	**11.2±0.5**	0.1
	AIS-Impute	**11.2±0.4**	0.1
	ADMM(scaled)	**11.2±0.7**	0.1
$m = 2000$	APG	**11.1±0.6**	0.5
	AIS-Impute(exact)	**11.1±0.6**	0.4
	AIS-Impute	**11.2±0.6**	0.4
	ADMM(scaled)	**11.3±0.8**	0.1

5.5 Multi-Relational Link Prediction

In this section, we perform experiments on the *YouTube* data set[11] [Lei *et al.*, 2009]. It contains 15,088 users, and describes five types of user interactions: contact, number of shared friends, number of shared subscriptions, number of shared subscribers, and the number of shared favorite videos. Thus, it forms a $15088 \times 15088 \times 5$ tensor, with a total of $27,257,790$ nonzero elements. Following [Chiang *et al.*, 2014], we formulate multi-relational link prediction as a tensor completion problem. As the observations are real-valued, we use the square loss in (23). Besides AIS-Impute (Algorithm 6), we also compare with the following state-of-the-art non-proximal-based tensor completion algorithms:

[11] http://socialcomputing.asu.edu/datasets/YouTube

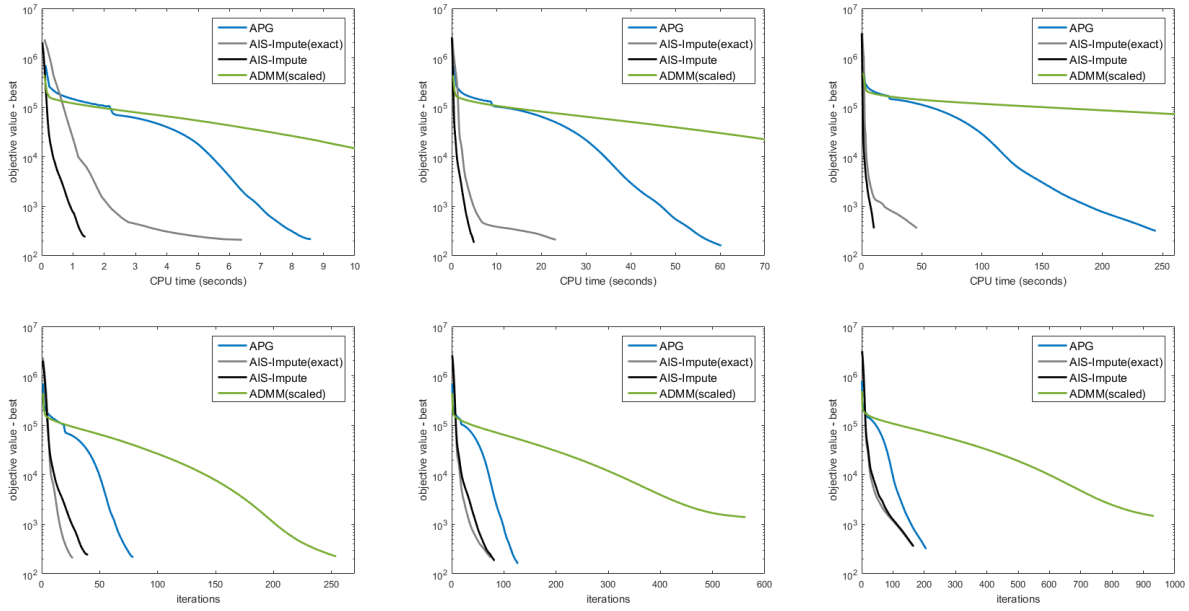

(a) $m = 500$. (b) $m = 1000$. (c) $m = 2000$.

Figure 5: Convergence of objective value on the synthetic tensor data. Top: vs CPU time (in seconds); Bottom: vs number of iterations.

1. GeomCG[12]: a gradient descent approach with gradients restricted on the Riemannian manifold [Kressner *et al.*, 2014];
2. An alternating direction method of multipliers approach (denoted "ADMM(overlap)") [Tomioka *et al.*, 2010][13], which solves the overlapping nuclear norm regularized tensor completion problem;
3. FaLRTC [Liu *et al.*, 2013][14]: It smooths the overlapping nuclear norm with Nesterov's method [Nesterov, 2005], and then solves the relaxed problem with accelerated gradient descent; and
4. TMac [Xu *et al.*, 2013][15]: An extension of LMaFit [Wen *et al.*, 2012] to tensor completion. It performs simultaneous low-rank matrix factorizations to all mode matricizations.

We do not compare with the ADMM(scaled), as it is much slower than AIS-Impute in Section 5.4.

YouTube Subset

First, we perform experiments on a small *YouTube* subset, obtained by random selecting 1000 users (leading to $12,101$ observations). We use 50% of the observations for training, another 25% for validation and the remaining for testing. Let $\mathcal{X}$ be the recovered tensor, and the testing ratings $\hat{O}_{ij}$ be indexed by the set $\hat{\Omega}$. For performance evaluation, we use (i) the testing root mean squared error RMSE $= \sqrt{\|P_{\hat{\Omega}}(\mathcal{X} - \hat{\mathcal{O}})\|_F^2 / \|\hat{\Omega}\|_1}$; and (ii) rank of the unfolded matrix in each mode. The experiments are repeated five times.

Performance is shown in Table 12 and Figure 6(a) shows the time comparison. ADMM(overlap) and FaLRTC have similar recovery performance, but are all very slow due to usage of SVD. As the overlapping nuclear norm is smoothed in FaLRTC, its cannot exactly recover a low-rank tensor. TMac is fast, but has the worst recovery performance. AIS-Impute enjoys fast speed and good recovery performance.

Table 12: Results on the *YouTube* subset. The lowest and comparable RMSEs (according to the pairwise t-test with 95% confidence) are highlighted.

	RMSE	rank of mode		
		1	2	3
GeomCG	0.672±0.050	7	7	5
ADMM(overlap)	0.690±0.030	142	142	5
FaLRTC	0.672±0.032	1000	1000	5
TMac	0.786±0.027	4	4	0
AIS-Impute	**0.616±0.029**	33	33	0

Full YouTube Data

Next, we perform experiments on the full *YouTube* data set with the same setup. As ADMM(overlap) and FaLRTC are too slow, we only compare GeomCG, TMac and AIS-Impute. Experiments are repeated five times.

[12] http://web.math.princeton.edu/~bartv/tensor_completion.html

[13] https://github.com/ryotat/tensor

[14] http://www.cs.rochester.edu/u/jliu/publications.html

[15] http://www.math.ucla.edu/~wotaoyin/papers/codes/TMac.zip

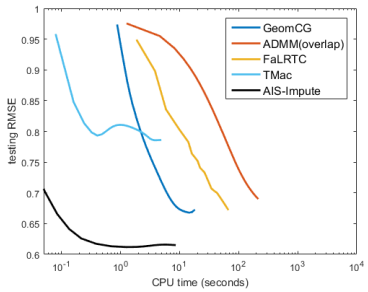

(a) data subset.

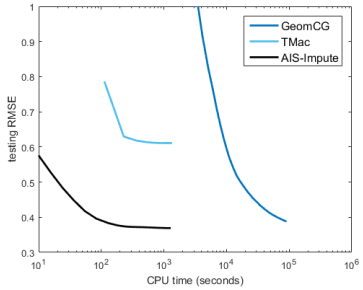

(b) full data set.

Figure 6: Testing RMSE vs CPU time on the *Youtube* data set.

Results are shown in Table 13, and Figure 6(b) shows the time. TMac has much worse performance than GeomCG and AIS-Impute. GeomCG is based on the (nonconvex) Turker decomposition, and its convergence rate is unknown. Moreover, its iteration time complexity has a worse dependency on the tensor rank than AIS-Impute ($\prod_{i=1}^{D} r_t^d$ vs $\sum_{i=1}^{D} r_t^d$), and thus GeomCG becomes very slow when the tensor rank is large. Overall, AIS-Impute has fast speed and good recovery performance.

Table 13: Results on the full *YouTube* dataset. The lowest and comparable RMSEs (according to the pairwise t-test with 95% confidence) are highlighted.

	RMSE	rank of mode		
		1	2	3
GeomCG	0.388±0.001	51	51	5
TMac	0.611±0.007	10	10	0
AIS-Impute	**0.369±0.006**	70	70	0

6 Conclusion

In this paper, we show that Soft-Impute, as a proximal algorithm, can be accelerated without losing the "sparse plus low-rank" structure crucial to the efficiency of Soft-Impute. To further reduce the per-iteration time complexity, we proposed an approximate-SVT scheme based on the power method. Theoretical analysis shows that the proposed algorithm still enjoys the fast $O(1/T^2)$ convergence rate. We also extend the proposed algorithm to handle low-rank tensor completion with the scaled latent nuclear norm as regularizer. We show that it is possible to preserve the "sparse plus low-rank" structure and fast $O(1/T^2)$ convergence. Extensive experiments on both synthetic and real-world data sets show that the proposed algorithm is much faster than the state-of-the-art.

As for future work, it will be interesting to develop parallel and distributed versions of the proposed algorithms for better speedup and scalability. Moreover, recommender systems usually have rich side information, such as the users' relationships, movies' content [Koren, 2008; Ma *et al.*, 2011; Zhao *et al.*, 2015]. It will be useful to incorporate these information into the algorithm for improved performance.

References

[Adomavicius and Tuzhilin, 2005] G. Adomavicius and A. Tuzhilin. Toward the next generation of recommender systems: A survey of the state-of-the-art and possible extensions. *IEEE Transactions on Knowledge and Data Engineering*, 17(6):734–749, 2005.

[Arbenz, 2010] P. Arbenz. Lecture notes on solving large scale eigenvalue problems, 2010. `http://people.inf.ethz.ch/arbenz/ewp/Lnotes/lsevp.pdf`.

[Avron *et al.*, 2012] H. Avron, S. Kale, V. Sindhwani, and S.P. Kasiviswanathan. Efficient and practical stochastic subgradient descent for nuclear norm regularization. In *Proceedings of the 29th International Conference on Machine Learning*, pages 1231–1238, 2012.

[Beck and Teboulle, 2009] A. Beck and M. Teboulle. A fast iterative shrinkage-thresholding algorithm for linear inverse problems. *SIAM Journal on Imaging Sciences*, 2(1):183–202, 2009.

[Boyd *et al.*, 2011] S.P. Boyd, N. Parikh, E. Chu, B. Peleato, and J. Eckstein. Distributed optimization and statistical learning via the alternating direction method of multipliers. *Foundations and Trends in Machine Learning*, pages 1–122, 2011.

[Cai *et al.*, 2010] J.-F. Cai, E.J. Candès, and Z. Shen. A singular value thresholding algorithm for matrix completion. *SIAM Journal on Optimization*, 20(4):1956–1982, 2010.

[Candès and Recht, 2009] E.J. Candès and B. Recht. Exact matrix completion via convex optimization. *Foundations of Computational Mathematics*, 9(6):717–772, 2009.

[Chiang *et al.*, 2014] K.-Y. Chiang, C.-J. Hsieh, N. Natarajan, I.S. Dhillon, and A. Tewari. Prediction and clustering in signed networks: A local to global perspective. *Journal of Machine Learning Research*, 15(1):1177–1213, 2014.

[Combettes and Wajs, 2005] P.L. Combettes and V.R. Wajs. Signal recovery by proximal forward-backward splitting. *Multiscale Modeling & Simulation*, 4(4):1168–1200, 2005.

[Frank and Wolfe, 1956] M. Frank and P. Wolfe. An algorithm for quadratic programming. *Naval Research Logistics*, 3(1-2):95–110, 1956.

[Haeffele *et al.*, 2014] Benjamin Haeffele, Eric Young, and Rene Vidal. Structured low-rank matrix factorization: Optimality, algorithm, and applications to image processing. In *Proceedings of the 31st International Conference on Machine Learning*, pages 2007–2015, 2014.

[Halko *et al.*, 2011] N. Halko, P.-G. Martinsson, and J.A. Tropp. Finding structure with randomness: Probabilistic algorithms for constructing approximate matrix decompositions. *SIAM Review*, 53(2):217–288, 2011.

[Hastie *et al.*, 2015] T. Hastie, R. Mazumder, J.D. Lee, and R. Zadeh. Matrix completion and low-rank SVD via fast alternating least squares. *Journal of Machine Learning Research*, 16:3367–3402, 2015.

[He and Yuan, 2012] B. He and X. Yuan. On the $O(1/n)$ convergence rate of the douglas-rachford alternating direction method. *SIAM Journal on Numerical Analysis*, 50(2):700–709, 2012.

[Hsieh and Olsen, 2014] C.-J. Hsieh and P. Olsen. Nuclear norm minimization via active subspace selection. In *Proceedings of the 31st International Conference on Machine Learning*, pages 575–583, 2014.

[Jacob *et al.*, 2009] L. Jacob, G. Obozinski, and J.-P. Vert. Group lasso with overlap and graph lasso. In *Proceedings of the 26th International Conference on Machine Learning*, pages 433–440, 2009.

[Ji and Ye, 2009] S. Ji and J. Ye. An accelerated gradient method for trace norm minimization. In *Proceedings of the 26th International Conference on Machine Learning*, pages 457–464, 2009.

[Kazienko *et al.*, 2011] P. Kazienko, K. Musial, and T. Kajdanowicz. Multidimensional social network in the social recommender system. *IEEE Transactions on Systems, Man, and Cybernetics*, 4(41):746–759, 2011.

[Kim and Leskovec, 2011] M. Kim and J. Leskovec. The network completion problem: inferring missing nodes and edges in networks. In *Proceedings of the 11st International Conference on Data Mining*, pages 47–58, 2011.

[Kolda and Bader, 2009] T.G. Kolda and B.W. Bader. Tensor decompositions and applications. *SIAM Review*, 51(3):455–500, 2009.

[Koren, 2008] Y. Koren. Factorization meets the neighborhood: a multifaceted collaborative filtering model. In *Proceedings of the 14th International Conference on Knowledge Discovery and Data Mining*, pages 426–434, 2008.

[Kressner *et al.*, 2014] D. Kressner, M. Steinlechner, and B. Vandereycken. Low-rank tensor completion by Riemannian optimization. *BIT Numerical Mathematics*, 54(2):447–468, 2014.

[Larsen, 1998] R.M. Larsen. Lanczos bidiagonalization with partial reorthogonalization. DAIMI PB-357, Department of Computer Science, Aarhus University, 1998.

[Lei *et al.*, 2009] T. Lei, X. Wang, and H. Liu. Uncoverning groups via heterogeneous interaction analysis. In *IEEE International Conference on Data Mining*, pages 503–512, 2009.

[Liu *et al.*, 2013] J. Liu, P. Musialski, P. Wonka, and J. Ye. Tensor completion for estimating missing values in visual data. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 35(1):208–220, 2013.

[Ma *et al.*, 2011] H. Ma, D. Zhou, C. Liu, M.R. Lyu, and I. King. Recommender systems with social regularization. In *Proceedings of the 4th International Conference on Web Search and Data Mining*, pages 287–296, 2011.

[Mazumder *et al.*, 2010] R. Mazumder, T. Hastie, and R. Tibshirani. Spectral regularization algorithms for learning large incomplete matrices. *Journal of Machine Learning Research*, 11:2287–2322, 2010.

[Mishra *et al.*, 2013] B. Mishra, G. Meyer, F. Bach, and R. Sepulchre. Low-rank optimization with trace norm penalty. *SIAM Journal on Optimization*, 23(4):2124–2149, 2013.

[Nesterov, 2005] Y. Nesterov. Smooth minimization of non-smooth functions. *Mathematical Programming*, 103(1):127–152, 2005.

[Nesterov, 2013] Y. Nesterov. Gradient methods for minimizing composite functions. *Mathematical Programming*, 140(1):125–161, 2013.

[Nocedal and Wright, 1999] J. Nocedal and S.J. Wright. *Numerical optimization*. Springer, 1999.

[O'Donoghue and Candès, 2012] B. O'Donoghue and E.J. Candès. Adaptive restart for accelerated gradient schemes. *Foundations of Computational Mathematics*, pages 1–18, 2012.

[Parikh and Boyd, 2014] N. Parikh and S. Boyd. Proximal algorithms. *Foundations and Trends in Optimization*, 1(3):127–239, 2014.

[Recht *et al.*, 2010] B. Recht, M. Fazel, and P.A. Parrilo. Guaranteed minimum-rank solutions of linear matrix equations via nuclear norm minimization. *SIAM Review*, 52(3):471–501, 2010.

[Schmidt *et al.*, 2011] M. Schmidt, N.L. Roux, and F.R. Bach. Convergence rates of inexact proximal-gradient methods for convex optimization. In *Advances in Neural Information Processing Systems*, pages 1458–1466, 2011.

[Shin *et al.*, 2017] K. Shin, L. Sael, and U. Kang. Fully scalable methods for distributed tensor factorization. *IEEE Transactions on Knowledge and Data Engineering*, 29(1):100–113, 2017.

[Tang *et al.*, 2008] L. Tang, H. Liu, J. Zhang, and Z. Nazeri. Community evolution in dynamic multi-mode networks. In *Proceedings of the 14th International Conference on Knowledge Discovery and Data Mining*, pages 677–685, 2008.

[Tibshirani, 2010] R. Tibshirani. Proximal gradient descent and acceleration. Lecture Notes, 2010. `http://www.stat.cmu.edu/~ryantibs/convexopt-S15/lectures/08-prox-grad.pdf`.

[Toh and Yun, 2010] K.-C. Toh and S. Yun. An accelerated proximal gradient algorithm for nuclear norm regularized linear least squares problems. *Pacific Journal of Optimization*, 6(615-640):15, 2010.

[Tomioka *et al.*, 2010] R. Tomioka, K. Hayashi, and H. Kashima. Estimation of low-rank tensors via convex optimization. Technical Report arXiv:1010.0789, Department of Mathematical Informatics, University of Tokyo, 2010.

[Wen *et al.*, 2012] Z. Wen, W. Yin, and Y. Zhang. Solving a low-rank factorization model for matrix completion by a nonlinear successive over-relaxation algorithm. *Mathematical Programming Computation*, 4(4):333–361, 2012.

[Wimalawarne *et al.*, 2014] K. Wimalawarne, M. Sugiyama, and R. Tomioka. Multitask learning meets tensor factorization: task imputation via convex optimization. In *Advances in Neural Information Processing Systems*, pages 2825–2833, 2014.

[Wu and Simon, 2000] K. Wu and H. Simon. Thick-restart Lanczos method for large symmetric eigenvalue problems. *SIAM Journal on Matrix Analysis and Applications*, 22(2):602–616, 2000.

[Xu *et al.*, 2013] Y. Xu, R. Hao, W. Yin, and Z. Su. Parallel matrix factorization for low-rank tensor completion. *Inverse Problems & Imaging*, 9(2), 2013.

[Yao and Kwok, 2015a] Q. Yao and J. T. Kwok. Accelerated inexact soft-impute for fast large-scale matrix completion. In *Proceedings of the 24th International Joint Conference on Artificial Intelligence*, pages 4002–4008, 2015.

[Yao and Kwok, 2015b] Q. Yao and J. T. Kwok. Colorization by patch-based local low-rank matrix completion. In *Proceedings of the 29th AAAI Conference on Artificial Intelligence*, pages 1959–1965, 2015.

[Zhang *et al.*, 2012] X. Zhang, D. Schuurmans, and Y.-L. Yu. Accelerated training for matrix-norm regularization: A boosting approach. In *Advances in Neural Information Processing Systems*, pages 2906–2914, 2012.

[Zhao *et al.*, 2015] Z. Zhao, L. Zhang, X. He, and W. Ng. Expert finding for question answering via graph regularized matrix completion. *IEEE Transactions on Knowledge and Data Engineering*, 27(4):993–1004, 2015.

A Proofs

A.1 Proposition 3.1

Proof. For any $X, Y \in \mathbb{R}^{m\times n}$,

$$\begin{aligned}
&\|\nabla f(X) - \nabla f(Y)\|_F^2 \\
&= \sum_{(i,j)\in\Omega} \left[\frac{d\ell(X_{ij}, O_{ij})}{dX_{ij}} - \frac{d\ell(Y_{ij}, O_{ij})}{dY_{ij}}\right]^2 \\
&\le \sum_{(i,j)\in\Omega} \rho^2 (X_{ij} - Y_{ij})^2 \qquad (30)\\
&\le \rho^2 \|X - Y\|_F^2,
\end{aligned}$$

where (30) follows from the fact that ℓ being ρ-Lipschitz smooth. Thus, $f(X)$ is ρ-Lipschitz smooth. □

A.2 Proposition 3.3

Proof. Let the SVD of $\breve{Z}_t$ be $U\Sigma V^\top$, and

$$\breve{Z}_t = [U_{k_t}; U_\perp] \begin{bmatrix}\Sigma_{k_t} & \\ & \Sigma_\perp\end{bmatrix} [V_{k_t}; V_\perp]^\top, \qquad (31)$$

where U_{k_t} contains the top k_t columns of U, and $U_\perp$ contains the remaining columns (and similarly for $(\Sigma_{k_t}, \Sigma_\perp)$ in Σ, and $(V_{k_t}, V_\perp)$ in V).

Lemma A.1. *The SVD of $Q^\top \breve{Z}_t$ is $(Q^\top U_{k_t})\Sigma_{k_t} V_{k_t}^\top$.*

Proof. As $Q^\top U_\perp = 0$,

$$\begin{aligned}
&Q^\top \breve{Z}_t \\
&= [Q^\top U_{k_t}; Q^\top U_\perp] \begin{bmatrix}\Sigma_{k_t} & \\ & \Sigma_\perp\end{bmatrix} [V_{k_t}; V_\perp]^\top \\
&= [Q^\top U_{k_t}; 0] \begin{bmatrix}\Sigma_{k_t} & \\ & \Sigma_\perp\end{bmatrix} [V_{k_t}; V_\perp]^\top = (Q^\top U_{k_t})\Sigma_{k_t} V_{k_t}^\top.
\end{aligned}$$

Moreover, as $\text{span}(Q) = \text{span}(U_{k_t})$, we have

$$(Q^\top U_{k_t})^\top (Q^\top U_{k_t}) = U_{k_t}^\top Q Q^\top U_{k_t} = U_{k_t}^\top U_{k_t} U_{k_t}^\top U_{k_t} = I.$$

Thus, $Q^\top U_{k_t}$ is orthogonal. □

From Lemma 2.2 and (31),

$$\begin{aligned}
&\text{SVT}_\lambda(\breve{Z}_t) \\
&= [U_{k_t}; U_\perp] \left(\begin{bmatrix}\Sigma_{k_t} & \\ & \Sigma_\perp\end{bmatrix} - \lambda I \right)_+ [V_{k_t}; V_\perp]^\top \\
&= U_{k_t} (\Sigma_{k_t} - \lambda I)_+ V_{k_t}^\top \qquad (32)\\
&= \left[U_{k_t} U_{k_t}^\top\right] U_{k_t} (\Sigma_{k_t} - \lambda I)_+ V_{k_t}^\top \\
&= Q \left[Q^\top U_{k_t} (\Sigma_{k_t} - \lambda I)_+ V_{k_t}^\top\right] \qquad (33)\\
&= Q\ \text{SVT}_\lambda(Q^\top \breve{Z}_t). \qquad (34)
\end{aligned}$$

Here, (32) follows from $k_t \ge \breve{k}_t$ and that $\breve{k}_t$ singular values of $\breve{Z}_t$ are larger than λ; (33) follows from $\text{span}(Q) = \text{span}(U_{k_t})$; and (34) follows from Lemma A.1. □

A.3 Proposition 3.4

Before proof of Proposition 3.4, we first introduce some Lemmas (Lemma A.2, A.4, A.3 and A.7) and Propositions (Proposition A.5 and A.6).

Lemma A.2 ([Combettes and Wajs, 2005])**.** *For any matrices A and B, $\|SVT_\lambda(A) - SVT_\lambda(B)\|_F \le \|A - B\|_F$.*

Let $Z_t^* \equiv \text{SVT}_{\mu l}(\breve{Z}_t)$, $\beta_t \equiv \|\breve{Z}_t\|_F$ and $\eta_t = \frac{\sigma_{k+1}(\breve{Z}_t)}{\sigma_k(\breve{Z}_t)}$.

Lemma A.3 ([Arbenz, 2010])**.** *Let the input to Algorithm 3 be $\breve{Z}_t$, and its top k left singular vectors be contained in U_k. Then, for $j = 0, 1, 2, \ldots$,*

$$\|Q_j Q_j{}^\top - U_k U_k^\top\|_F \le \eta_t^j \alpha_t,$$

where $\alpha_t = \|Q_0 Q_0{}^\top - U_k U_k^\top\|_F$ and Q_0 is the span of $\breve{Z}_t R_t$.

Lemma A.4. *For output $\tilde{X} = (QU)\Sigma V^\top$ from Algorithm 4, we have $\|\tilde{X} - Z_t^*\|_F \le \|U_k U_k^\top - QQ^\top\|_F \beta_t$.*

Proof. From Proposition 3.3,

$$\begin{aligned} Z_t^* - \tilde{X} &= \text{SVT}_{\mu\lambda}(\breve{Z}_t) - Q\text{SVT}_{\mu\lambda}(Q^\top \breve{Z}_t) \\ &= \text{SVT}_{\mu\lambda}(U_k U_k^\top \breve{Z}_t) - \text{SVT}_{\mu\lambda}(QQ^\top \breve{Z}_t). \end{aligned}$$

Using Lemma A.2 and the Cauchy's inequality,

$$\begin{aligned} \|\tilde{X} - Z_*\|_F &= \|\text{SVT}_{\mu\lambda}(U_k U_k^\top \breve{Z}_t) - \text{SVT}_{\mu\lambda}(QQ^\top \breve{Z}_t)\|_F \\ &\le \|(U_k U_k^\top - QQ^\top)\breve{Z}_t\|_F \\ &\le \|U_k U_k^\top - QQ^\top\|_F \beta_t, \end{aligned}$$

and result follows. □

Proposition A.5. *Let $G_t \in \partial h_{\mu\lambda\|\cdot\|_*}(\tilde{X}; \breve{Z}_t)$, then $\|G_t\|_F$ is upper-bounded by a constant γ_t.*

Proof. Let the reduced SVD of $\tilde{X}$ be $U\Sigma V^\top$ (only positive singular values are contained). By the definition of subgradient of the nuclear norm [Candès and Recht, 2009],

$$\partial h_{\mu\lambda\|\cdot\|_*}(\tilde{X}; \breve{Z}_t) = \tilde{X} - \breve{Z}_t + \mu\lambda(UV^\top + W),$$

where

$$W^\top U = 0,\ WV = 0,\ \text{and } \|W\|_\infty \le 1. \tag{35}$$

Thus,

$$\begin{aligned} \|G_t\|_F &= \|\tilde{X} - \breve{Z}_t + \mu\lambda(UV^\top + W)\|_F \\ &\le \|\tilde{X} - \breve{Z}_t\|_F + \mu\lambda\|UV^\top + W\|_F. \end{aligned} \tag{36}$$

For the first term in (36),

$$\begin{aligned} &\|\tilde{X} - \breve{Z}_t\|_F \\ &= \|\tilde{X} - Z_t^* + Z_t^* - \breve{Z}_t\|_F \\ &\le \|\tilde{X} - Z_t^*\|_F + \|Z_t^* - \breve{Z}_t\|_F \\ &= \|Z_t^* - \breve{Z}_t\|_F + \|Z_t^* - Q\text{SVT}_{\mu\lambda}(Q^\top \breve{Z}_t)\|_F \\ &= \|Z_t^* - \breve{Z}_t\|_F + \|Z_t^* - \tilde{X}\|_F \\ &\le \|Z_t^* - \breve{Z}_t\|_F + \|U_k U_k^\top - QQ^\top\|_F \beta_t \quad (37) \\ &\le \|Z_t^* - \breve{Z}_t\|_F + \alpha_t\beta_t. \quad (38) \end{aligned}$$

Here, (37) follows from Lemma A.4, and (38) from Lemma A.3. As $\|W\|_\infty \le 1$ from (35), thus

$$\|W\|_F = \sqrt{\sum_{i=1}^m \sigma_i^2(W)} \le \sqrt{m}.$$

For the second term in (36), then

$$\begin{aligned} \|UV^\top + W\|_F &\le \sqrt{\text{tr}(U^\top UV^\top V)} + \|W\|_F \\ &\le \sqrt{k_t} + \sqrt{m} \le 2\sqrt{m}. \end{aligned} \tag{39}$$

Combining (38) and (39), by Lemma A.3:

$$\|G_t\|_F \le 2\mu\lambda\sqrt{m} + \|Z_t^* - \breve{Z}_t\|_F + \alpha_t\beta_t. \tag{40}$$

Since Z_t^* is independent of $\tilde{X}$, $\|Z_t^* - \breve{Z}_t\|_F$ is a constant. Hence, $\|G_t\|_F$ is upper bounded by

$$\gamma_t = 2\mu\lambda\sqrt{m} + \|Z_t^* - \breve{Z}_t\|_F + \alpha_t\beta_t,$$

which a constant. □

Proposition A.6. *Assume that $k_t \ge \breve{k}_t$. Let $h_{\mu\lambda\|\cdot\|_*}(\tilde{X}; \breve{Z}_t)$ be as defined in* (8)*. Then, for Algorithm 4, we have*

$$h_{\mu\lambda\|\cdot\|_*}(\tilde{X}; \breve{Z}_t) \le h_{\mu\lambda\|\cdot\|_*}(Z_t^*; \breve{Z}_t) + \alpha_t\beta_t\gamma_t\eta_t^J.$$

Proof. As h is convex,

$$h_{\mu\lambda\|\cdot\|_*}(\tilde{X}; \breve{Z}_t) \le h_{\mu\lambda\|\cdot\|_*}(Z_t^*; \breve{Z}_t) + \text{tr}((\tilde{X} - Z_t^*)^\top G_t) \tag{41}$$

where $G_t \in \partial h_{\mu\lambda\|\cdot\|_*}(\tilde{X}; \breve{Z}_t)$. Next, we bound the second term on the r.h.s. of (41).

$$\begin{aligned} \text{tr}((\tilde{X} - Z_t^*)^\top G_t) &\le \|\tilde{X} - Z_t^*\|_F \|G_t\|_F \\ &\le \gamma_t \|\tilde{X} - Z_t^*\|_F \quad (42) \\ &\le \gamma_t\beta_t \|QQ^\top - U_k U_k^\top\|_F \quad (43) \\ &\le \eta_t^J(\alpha_t\beta_t\gamma_t). \quad (44) \end{aligned}$$

Here, (42) follows from Proposition A.5; (43) from Lemma A.4; and (44) from Lemma A.3. Result follows on combining (41) and (44). □

Lemma A.7. *If $\{F(X_t)\}$ is upper-bounded where F is the objective at* (19)*, then $\|X_t\|_F$ from Algorithm 5 is upper-bounded.*

Proof. As $\{F(X_t)\}$ is upper bounded and note that

$$F(X) \to +\infty \Leftrightarrow \|X\|_F \to +\infty.$$

for (19), then $\{\|X_t\|_F\}$ is also upper bounded. □

Now, we are ready to prove Proposition 3.4. As α_t, β_t and γ_t only depend on X_t, from Lemma A.7, they are all upper bounded. Let $q = \sup_t \alpha_t\beta_t\gamma_t$, and $q < \infty$ is a constant. Then by Proposition A.6, and note that Algorithm 4 is run for t iterations at tth loop of Algorithm 5. Let $\eta = \max_t \eta_t \in (0, 1)$, we have

$$h_{\mu\lambda\|\cdot\|_*}(X_{t+1}; \breve{Z}_t) \le h_{\mu\lambda\|\cdot\|_*}(Z_t^*; \breve{Z}_t) + \varepsilon_t.$$

Hence, $\varepsilon_t = q\eta^t$ decays at a linear rate.

A.4 Theorem 3.5

Proof. From Proposition 3.4, ε_t decays at a linear rate. Moreover, there is no error on the computation of gradient. Thus, conditions in Proposition 2.1 are satisfied, and Algorithm 5 converges with a rate of $O(1/T^2)$. □

A.5 Proposition 4.1

Proof. Note that

$$\begin{aligned} &\min_{\mathfrak{X}^1,\ldots,\mathfrak{X}^D} \frac{1}{2}\|[\mathfrak{X}^1,\ldots,\mathfrak{X}^D] - [\breve{\mathfrak{Z}}_t^1,\ldots,\breve{\mathfrak{Z}}_t^D]\|_F^2 + \mu\sum_{d=1}^D \lambda_d \|\mathfrak{X}_{\langle d\rangle}^d\|_* \\ &= \sum_{d=1}^D \min_{\mathfrak{X}^d} \frac{1}{2}\|\mathfrak{X}^d - \breve{\mathfrak{Z}}_t^d\|_F^2 + \mu\lambda_d\|\mathfrak{X}_{\langle d\rangle}^d\|_*, \\ &= \sum_{d=1}^D \min_{\mathfrak{X}^d} \frac{1}{2}\|\mathfrak{X}_{\langle d\rangle}^d - (\breve{\mathfrak{Z}}_t^d)_{\langle d\rangle}\|_F^2 + \mu\lambda_d\|\mathfrak{X}_{\langle d\rangle}^d\|_*. \end{aligned} \tag{45}$$

The $\mathfrak{X}^d$'s in (45) are independent of each other, and

$$(\text{SVT}_{\mu\lambda_d}(\breve{\mathfrak{Z}}_{\langle d\rangle}^d))_{\langle d\rangle} = \arg\min_{\mathfrak{X}^d} \frac{1}{2}\|\mathfrak{X}_{\langle d\rangle}^d - \breve{\mathfrak{Z}}_{\langle d\rangle}^d\|_F^2 + \mu\lambda_d\|\mathfrak{X}_{\langle d\rangle}^d\|_*.$$

and thus result follows. □

A.6 Proposition 4.2

Proof. For any $\mathcal{X}^1, \ldots, \mathcal{X}^D$, $\mathcal{Y}^1, \cdots, \mathcal{Y}^D$, and let $\tilde{\mathcal{X}} = \sum_{d=1}^D \mathcal{X}^d$ and $\tilde{\mathcal{Y}} = \sum_{d=1}^D \mathcal{Y}^d$.

$$\begin{aligned}
&\|\nabla f([\mathcal{X}^1, \ldots, \mathcal{X}^D]) - \nabla f([\mathcal{Y}^1, \cdots, \mathcal{Y}^D])\|_F^2 \\
&= \sum_{(i_1,\ldots,i_D)\in\Omega} \left[\frac{d\ell(\tilde{\mathcal{X}}_{i_1\ldots i_D}, \mathcal{O}_{i_1\ldots i_D})}{d\tilde{\mathcal{X}}_{i_1\ldots i_D}} - \frac{d\ell(\tilde{\mathcal{Y}}_{i_1\ldots i_D}, \mathcal{O}_{i_1\ldots i_D})}{d\tilde{\mathcal{Y}}_{i_1\ldots i_D}} \right]^2 \\
&\leq \sum_{(i_1,\ldots,i_D)\in\Omega} \rho^2 \left(\tilde{\mathcal{X}}_{i_1\ldots i_D} - \tilde{\mathcal{Y}}_{i_1\ldots i_D} \right)^2 \leq \rho^2 \left\| \tilde{\mathcal{X}} - \tilde{\mathcal{Y}} \right\|_F^2,
\end{aligned}$$

where the first inequality comes from the ρ-Lipschitz smoothness of ℓ. Note that

$$\begin{aligned}
\left\| \tilde{\mathcal{X}} - \tilde{\mathcal{Y}} \right\|_F^2 &\leq D \sum_{d=1}^D \|\mathcal{X}^d - \mathcal{Y}^d\|_F^2 \\
&= D\|[\mathcal{X}^1, \ldots, \mathcal{X}^D] - [\mathcal{Y}^1, \ldots, \mathcal{Y}^D]\|_F^2.
\end{aligned}$$

We have

$$\begin{aligned}
&\|\nabla f([\mathcal{X}^1, \ldots, \mathcal{X}^D]) - \nabla f([\mathcal{Y}^1, \ldots, \mathcal{Y}^D])\|_F \\
&\leq \sqrt{D}\rho \|[\mathcal{X}^1, \ldots, \mathcal{X}^D] - [\mathcal{Y}^1, \ldots, \mathcal{Y}^D]\|_F,
\end{aligned}$$

and thus f is $\sqrt{D}\rho$-Lipschitz smooth. □

A.7 Theorem 4.3

Proof. From the definition of h in (8),

$$\begin{aligned}
&h_{\mu g}\left([\mathcal{X}_{t+1}^1, \ldots, \mathcal{X}_{t+1}^D]; [\breve{\mathcal{Z}}_t^1, \ldots, \breve{\mathcal{Z}}_t^D] \right) \\
&= \sum_{d=1}^D \frac{1}{2} \left\| (\mathcal{X}_{t+1}^d)_{\langle d \rangle} - (\breve{\mathcal{Z}}_t)_{\langle d \rangle}^d \right\|_F^2 + \mu\lambda_d \|(\mathcal{X}_{t+1}^d)_{\langle d \rangle}\|_*, \\
&= \sum_{d=1}^D h_{\mu\lambda_d\|\cdot\|_*} \left((\mathcal{X}_{t+1}^d)_{\langle d \rangle}; (\breve{\mathcal{Z}}_t)_{\langle d \rangle}^d \right). \qquad (46)
\end{aligned}$$

As proximal step is inexact in Algorithm 6, using Proposition A.6 on (46),

$$\begin{aligned}
&h_{\mu\lambda_d\|\cdot\|_*} \left((\mathcal{X}_{t+1}^d)_{\langle d \rangle}; (\breve{\mathcal{Z}}_t)_{\langle d \rangle}^d \right) \\
&\leq h_{\mu\lambda_d\|\cdot\|_*} \left((\mathcal{W}_*^d)_{\langle d \rangle}; (\breve{\mathcal{Z}}_t)_{\langle d \rangle}^d \right) + (\alpha_d)_t (\beta_d)_t (\gamma_d)_t (\eta_d)_t^J,
\end{aligned}$$

where $(\mathcal{W}_*^d)_{\langle d \rangle} = \text{SVT}_{\mu\lambda_d\|\cdot\|_*} \left((\breve{\mathcal{Z}}_t)_{\langle d \rangle}^d \right)$, and α_d, β_d, γ_d, η_d are constants depending on $(\breve{\mathcal{Z}}_t)_{\langle d \rangle}^d$. Let $(c_d)_t = (\alpha_d)_t (\beta_d)_t (\gamma_d)_t$. As $J = t$,

$$\begin{aligned}
&h_{\mu g}\left([\mathcal{X}_{t+1}^1, \cdots, \mathcal{X}_{t+1}^D]; [\breve{\mathcal{Z}}_t^1, \cdots, \breve{\mathcal{Z}}_t^D] \right) \qquad (47) \\
&\leq h_{\mu g}\left([\mathcal{W}_*^1, \cdots, \mathcal{W}_*^D]; [\breve{\mathcal{Z}}_t^1, \cdots, \breve{\mathcal{Z}}_t^D] \right) + \sum_{d=1}^D (c_d)_t (\eta_d)_t^t.
\end{aligned}$$

As $F([\mathcal{X}_t^1, \ldots, \mathcal{X}_t^D])$ is upper-bounded and

$$\lim_{\|\mathcal{X}^d\|_F \to \infty} F([\mathcal{X}_t^1, \ldots, \mathcal{X}_t^D]) = \infty,$$

for any $d = 1, \ldots, D$. Then, $\|\mathcal{X}_t^d\|_F$ for $d = 1, \ldots, D$ are also upper-bounded. Thus,

$$q = \sup_t \sum_{d=1}^D (c_d)_t < \infty.$$

Let $\eta = \max_{d,t}((\eta_d)_t) < 1$. Together with (47), we have

$$\begin{aligned}
&h_{\mu g}\left([\mathcal{X}_{t+1}^1, \ldots, \mathcal{X}_{t+1}^D]; [\breve{\mathcal{Z}}_t^1, \ldots, \breve{\mathcal{Z}}_t^D] \right) \\
&\leq h_{\mu g}\left([\mathcal{W}_*^1, \ldots, \mathcal{W}_*^D]; [\breve{\mathcal{Z}}_t^1, \ldots, \breve{\mathcal{Z}}_t^D] \right) + \varepsilon_t,
\end{aligned}$$

and the approximation error $\varepsilon_t = q\eta^t$ decays at a linear rate. Moreover, there is no error on the computation of gradient. Thus, the conditions in Proposition 2.1 are satisfied, and Algorithm 6 converges with a rate of $O(1/T^2)$. □